

INSIGHT MEDIA

Projects

2026

MANUFACTURING TRANSFORMED

THE RESEARCH, TECHNOLOGY AND PARTNERSHIPS BRINGING INDUSTRIAL INNOVATION TO THE FACTORY FLOOR

EFFRA

MEDIA PARTNERS WITH EFFRA'S MANUFACTURING PARTNERSHIP DAYS 2026



Even great research won't speak for itself.

So we help it find its voice.

Get your **story heard.**
Get your **audiences engaged.**

Projects Magazine
Telling your story to the wider research and industry community

Insight Access
Flexible communications, dissemination and engagement support when you need it

Project partnerships
Shaping focused CD&E plans that make reviewers take notice

INSIGHT MEDIA

Connecting research with people

Media partner for EFFRA's Manufacturing Partnership Days 2026

Projects
Showcasing research and innovation in Europe

EDITORIAL
Editorial Director William Davis
wdavis@insightm.co.uk
Editor David Hatvani
info@insightm.co.uk
Editorial Assistant Rhianna Mone
info@insightm.co.uk

PRODUCTION
Production Manager Lydia Faircloth
lydia@insightm.co.uk
Graphic Designer Doms Valentin
dom@insightm.co.uk
Client Services Manager Steve Harper
info@insightm.co.uk

PUBLISHING
Managing Director William Davis
wdavis@insightm.co.uk
Director, Communications Sam Davis
sam@insightm.co.uk
Dissemination Consultant Jarrod McTeggart
Jarrod@insightm.co.uk
Social Media Manager Daryl Vicentepotazo
daryl@insightm.co.uk
Web Developer Greg Kett
greg@insightm.co.uk

INSIGHT MEDIA

Insight Media Group Ltd

12/13 King Square,
Bristol BS2 8JH,
United Kingdom

t: +44 (0)117 248 1755

e: sam@insightm.co.uk
www.insightm.co.uk

© Insight Media Group Ltd November 2026
ISSN 2040-7335

Foreword

The future of European manufacturing depends on its ability to innovate, adapt and make better use of its resources, skills and expertise. Faced with intensifying global competition, pressure on costs and resources, and the need for more resilient, greener production, Europe must now draw even more on its research strengths to help turn all that potential into lasting improvements on the factory floor.

The projects featured in this very special manufacturing edition of Projects Magazine demonstrate the breadth of work advancing that ambition. From artificial intelligence, digital twins and robotics to advanced materials, smarter monitoring and circular production, they approach manufacturing challenges from different directions. Their shared purpose is practical: more efficient processes, fewer defects, lower costs, less waste, and products and services that deliver greater value. Together, they show how technological progress can strengthen both individual businesses and the wider industrial economy.

Such work needs sustained public commitment, and the European Commission's funding programmes and calls provide just that. With essential support for collaborative research, it enables organisations to share expertise, test ideas and tackle challenges beyond the reach of any single business. Organisations such as EFFRA help connect these efforts with industrial priorities. We are delighted to share this edition at EFFRA's Manufacturing Partnership Days 2026, where those connections can develop further.

Yet a project's closing date rarely marks the end of the journey, with industrial adoption rarely following automatically on. Exploitation must accommodate the realities of manufacturing: investment decisions, validation, integration, skills and the time needed to build confidence in unfamiliar technologies. Industry must help shape that journey from the outset, while support mechanisms need to recognise that commercial readiness and funding deadlines do not always coincide. Continued access to expertise, demonstration facilities and partnerships, as well as funding, can give promising results the opportunity to mature.

At Insight Media, we see communication and engagement as crucial elements of this work. Giving great research its voice means helping people understand what has been achieved, why it matters and how they might use it. It also means creating opportunities for dialogue, so researchers hear industry's needs and potential users can engage with emerging solutions.

Our role is to help those connections take root and grow beyond publication and project completion. We hope these pages prompt conversations and the practical next steps that give excellent research the chance to flourish at industrial scale. If it can do that, then we are well on the way to having a more competitive, resource-efficient and resilient European manufacturing sector.

William Davis
Managing Director, Insight Media

Contents

NEWS AND FEATURES

6 **News**
The latest developments in manufacturing research, industrial innovation and European funding opportunities

9 **Keynote interview**
EFFRA's Executive Director Željko Pazin discusses how the new EFFRA Roadmap will help European manufacturing research connect with industrial priorities

24 **Made in Europe and the UK**
Why European manufacturing policy must recognise Britain's contribution to shared industrial success

31 **Zero-defect manufacturing**
Why Europe needs renewed investment to help turn promising research outcomes into lasting industrial impact.

58 **Insight Access**
Helping projects, research organisations and industry turn research into practical opportunities through communication, engagement and exploitation support

RESEARCH FOCUS: ZERO DEFECT MANUFACTURING

83 **FLASH-COMP**
Bringing sensing, predictive simulation and decision support together to Zero Defect Manufacturing in composite production

RESEARCH FOCUS: MATERIALS AND SUSTAINABLE PRODUCTION

12 **MAGELLAN**
Building a more resilient European supply chain through permanent magnet recycling and resource-efficient manufacturing

18 **SIMPLI-DEMO**
Using ultrasound and microwave technologies to make chemical production more efficient, flexible and sustainable

RESEARCH FOCUS: SMARTER AND MORE FLEXIBLE FACTORIES

26 **REED**
Connecting manufacturers through flexible production networks for large and complex components

32 **ONE4ALL**
Helping factories adapt through modular production systems and collaborative robotics

38 **AIRISE**
Supporting manufacturing SMEs in adopting artificial intelligence for more efficient and sustainable production

43 **SoliDAIR**
Bringing trustworthy AI, data and robotics into industrial inspection, quality control and process optimisation

RESEARCH FOCUS: ADVANCED SENSORS AND PHOTONICS

48 **LASTSTEP**
Developing silicon-compatible infrared lasers and detectors for a new generation of sensing applications

53 **SMARTSHAPE**
Combining smart polymers and microsensor technology to advance continuous blood pressure monitoring

RESEARCH FOCUS: CONNECTED INTELLIGENCE

60 **EXTRA-BRAIN**
Developing adaptable, explainable and energy-efficient AI inspired by the human brain

66 **OpenSwarm**
Enabling low-power smart devices to share information, coordinate decisions and work together

72 **MYRTUS**
Connecting devices, edge systems and cloud infrastructure through intelligent computing orchestration

78 **HYPER-AI**
Using self-organising computing networks to support demanding industrial applications

NEWS

AI prepares factory inspection before the first circuit board is built

Siemens and NVIDIA are developing manufacturing systems that prepare quality inspection before physical prototypes exist and help production planners respond to changing orders, bringing artificial intelligence into everyday factory decisions.

Demonstrated at Siemens' electronics factory in Erlangen, Germany, the applications combine digital twins, industrial software and AI to address two persistent manufacturing challenges: introducing new products quickly and adapting production without disrupting existing operations.

For printed circuit board assemblies, preparing automated optical inspection traditionally depends on having a physical prototype. Engineers must wait for sample boards before training vision systems and configuring inspection equipment, adding time between product design and production.

The new workflow uses product design, component and manufacturing information to create a digital twin of the assembly. NVIDIA Cosmos then helps generate realistic images representing both good boards and potential defects, allowing inspection preparation to begin before the first physical board reaches the line.

The collaboration also includes AI-supported secondary inspection, intended to help distinguish genuine defects from false alarms once production begins. Project targets include reducing manual engineering effort in inspection-programme generation by up to 85 per cent, alongside fewer false calls and improved first-pass yield.

A separate application supports production planning. When an order



arrives, an AI agent generates alternative scenarios that account for available materials, production capacity, delivery commitments and the costs of switching between products.

Planners can examine the consequences of different choices, while proposed changes are tested in a digital twin of the factory before implementation. A production controller reviews the options and approves the plan before permitted changes are applied through the connected planning system.

Siemens identifies potential cost savings of up to 20 per cent from the planning approach. The figures relate to specific applications and project targets rather than a reduction in total factory costs.

The Erlangen site manufactures power electronics and automation technology and serves as a working environment for developing and testing the systems during live production. The partners intend it to provide a blueprint for wider deployment.

For manufacturers handling frequent product changes, the work points towards shorter preparation cycles and more flexible scheduling, with simulation allowing teams to assess changes while physical production continues.

NEWS IN BRIEF

Robot vision

ABB Robotics has introduced AI Robot Trainer to help manufacturers develop and deploy robot vision applications. Combining synthetic and real images, the system supports tasks including picking, inspection, sorting and assembly, with a guided workflow intended to reduce reliance on specialist AI expertise.

Urban mining

The FutuRaM project, which closed in March 2026, has mapped secondary raw materials in European waste streams, including batteries, electronics and vehicles. Its Urban Mine Platform provides data and projections to 2050, helping manufacturers, recyclers and investors assess potential sources of recovered materials.

Bio-based fibres

Fraunhofer IAP is developing carbon fibres from renewable feedstocks including cellulose and lignin. By adjusting materials and processing conditions, researchers can tailor properties such as strength, porosity and electrical conductivity for potential applications in lightweight components, batteries and fuel cells.

Packaging partnership

The new PaperFlex Consortium is inviting businesses, innovators, suppliers and investors to support paper-based flexible packaging development. Founding members include Nestlé, Unilever and Mars. Planned collaborative trials will address coatings, product protection and the performance of new materials on manufacturing lines.

Smarter quality checks bring printed railway spare parts closer

Researchers in Germany have developed monitoring and inspection methods that could help manufacturers produce large plastic railway spare parts on demand, addressing a practical obstacle to wider industrial use of 3D printing.

Fraunhofer IMWS announced the results on the 15th of October as part of AddiQ, a collaborative project developing quality assurance across additive manufacturing. Its work focuses on fused granulate fabrication, which builds components using plastic granules rather than filament.

Railway vehicles can remain in service for decades, while individual components wear out or become unavailable. Producing small batches of replacement parts through conventional manufacturing can be expensive, creating an opportunity for additive processes.

However, printing a replacement is only part of the challenge. Manufacturers also need reliable ways to assess its properties and demonstrate that production can deliver consistent quality.

The wider AddiQ programme connects material testing, monitoring during printing and inspection of completed components. Its emphasis on individual parts and small production runs reflects applications where repeatability matters even when manufacturers are producing relatively few items.

Within the transport-focused work, SLV Halle is developing quality assurance and standardisation concepts for granulate-printed components. Drawing on experience from established metalworking rules, it aims to provide a framework for assessing and checking the properties of printed plastic parts.

Fraunhofer IMWS has complemented this work with sensors that record printing conditions and automated image analysis to assess nozzle wear. Linking those measurements to component tests is intended to make production deviations visible earlier and support better maintenance decisions.

The team demonstrated its approach using a 500-millimetre model of an ICE train front apron, approximately 38 per cent of the original component's size. The full-size part exceeded the printer's available build space.

The next step is to transfer the methods to customer components with industrial partners and qualify them for specific applications.

For manufacturers and transport operators, the work addresses a central question in on-demand production: how to combine the flexibility of printing replacement parts with the evidence needed to trust their quality.



NEWS IN BRIEF

Metal printing

Fraunhofer ILT has developed a fluidic switch that stops metal powder flow within milliseconds during laser material deposition, reducing waste during interruptions. The institute says the system operates without moving parts in the powder path and is already used in series production.

Power semiconductors

Polar Semiconductor and IC-Link by imec have announced an agreement to bring imec's gallium nitride technology to eight-inch wafer production in Minnesota. The collaboration targets power devices for data centres, automotive applications, aerospace and industrial electricity systems.

Recycling robotics

The EU-funded HARMONICA project, which started in June 2026, has outlined plans for semi-autonomous mobile robots and digital twins that support workers in hazardous recycling environments. Coordinated by Queen's University Belfast, the project will combine physical and cognitive worker models with information from machines and their surroundings.

Textile circularity

IGERS' new textile recycling plant in northern Italy combines automated sorting with material recovery and digital traceability. Technology partner STADLER says the facility can process four tonnes an hour, preparing recovered material from manufacturing and consumer waste for the spinning industry.

INSIGHT MEDIA

Insight Access

Insight Media's flexible communications partnership, giving research organisations on-demand, flexible access to Europe's leading experts in research communication, dissemination and impact.

Embed the experts into your team

Research communications

Dissemination & outreach

Stakeholder engagement

Impact & exploitation

Health & clinical communications

Content, digital & media production

Graphic design

How it works

Choose the level of monthly support that suits your organisation and access our experienced team whenever you need it.

Flexible monthly capacity

Predictable costs

Scale support up or down as priorities change

One trusted partner instead of multiple suppliers

Why choose Insight Access

No recruitment or onboarding

No procurement burden

Access to specialist expertise on demand

More capacity without increasing headcount

Trusted by research organisations across Europe for more than 15 years

Great research deserves outstanding communication

Book your introductory conversation today

Sam Davis
sam@insightm.co.uk
www.insightm.co.uk/insightaccess

INSIGHT FEATURE

EFFRA's roadmap for Europe's manufacturing future



As Europe prepares its next research and innovation funding programme, EFFRA has launched a draft roadmap that sets out a clear vision for competitive, resilient and sustainable manufacturing in Europe. Executive Director **Željko Pazin** explains why collaboration and support for industrial adoption are central to that vision, and key to Europe retaining the capabilities on which its security and prosperity depend.

When the European Factories of the Future Research Association (EFFRA) was established in 2009, manufacturing offered a route towards recovery from the financial crisis. Strengthening the companies that made Europe's products, and the technologies behind them, was central to rebuilding growth. Through the Factories of the Future partnership, EFFRA brought industry, research organisations and universities together to help shape European research priorities and develop production technologies industry could use. Today, EFFRA represents the private side of the Made in Europe manufacturing partnership under Horizon Europe.

The early emphasis was practical: better machines, faster processes, greater productivity and technologies that could improve environmental performance. EFFRA Executive Director Željko Pazin describes how that perspective subsequently widened: "We started looking beyond individual machines and the factory floor towards the company, logistics and the wider industrial ecosystem," he says. "The expectations of customers, employees and society became increasingly important."

That evolution explains the broad focus of EFFRA's new roadmap for the future of European manufacturing, launched earlier this year. The underlying ambition of this important piece of work remains competitive and sustainable European manufacturing. According to Željko, however, the circumstances in which that ambition must be achieved have changed considerably since EFFRA's formation in 2009.

European manufacturers face stronger competition from China across machinery and production technologies, putting pressure on capabilities that have long underpinned Europe's industrial strength. Improving productivity remains essential, but Europe must also ensure it retains the equipment suppliers, expertise and industrial networks needed to develop and manufacture future products.

At the same time, the pandemic, disruption to major shipping routes and geopolitical uncertainty have exposed the vulnerability of interconnected supply chains. Access to materials, components and technologies can no longer be taken for granted. Manufacturers need production systems that can adapt when supplies are interrupted, respond to changing demand and scale up quickly when new needs arise, including those associated with Europe's security.

These pressures have brought strategic autonomy and resilience to the forefront of EFFRA's thinking. The challenge now extends beyond just making an individual factory more efficient or sustainable to maintaining Europe's ability to produce, innovate and respond to events. That requires cooperation across companies, research organisations and supply networks – the broader industrial perspective reflected in the emerging roadmap.

The draft roadmap focuses on Europe's 2028–2034 funding period, and presents its proposals as five key pillars: (i) innovation, industrial scale-up and productivity; (ii) strategic autonomy and resilience; (iii) value retention and climate-neutral production; (iv) digital technologies and AI; and (v) human capital and human-centric manufacturing. Together, they describe the capabilities European industry will need and the conditions that will allow companies to deploy them.

The first pillar connects technological leadership with the ability to bring innovations into production at commercial scale. Advanced machinery, robotics, inspection and flexible manufacturing systems remain essential. There is also a growing emphasis on using AI to design and engineer production systems, alongside improving their operation. "As those systems become more complex, engineering itself can become a bottleneck," Željko adds. "Reducing that burden can shorten the journey from a new product idea to reliable production."

Alongside improving productivity, Željko sees resilience as crucial to Europe's manufacturing future. He believes companies need to be able to adapt production when circumstances change.

"Production systems that can be reconfigured quickly give businesses more options when materials become unavailable, demand changes or an emergency requires different products," he says.

The roadmap proposes a stronger understanding of critical dependencies, secure industrial infrastructure and greater capacity to move technologies between civilian and defence applications, so-called dual-use innovation. Željko sees great opportunities for Europe in this. "Manufacturing expertise developed in one sector can serve another, allowing existing businesses to enter new markets and helping Europe respond to changing security needs," he says. "The practical task is to make shared capabilities accessible and avoid treating every application as an entirely separate development effort."

The third pillar, value retention and climate-neutral production, focuses on how manufacturers can preserve the resources and value already invested in products, components and equipment. It advocates repair, reuse, refurbishment, remanufacturing and upgrading to keep them in productive use for longer. Recycling remains important, but it comes after opportunities to preserve more of the original product's function and value.

Making these activities work at industrial scale requires investment in inspection, disassembly and restoration technologies, supported by reliable information about a product's condition and history. Manufacturers also need ways to collect used products, certify their performance after repair or upgrading, and find customers for them. The roadmap therefore links technological development with the supply networks and commercial arrangements needed to make value retention a viable business.

"We need to create circular business models so that there is a motivation for companies to do it themselves," Željko says. "Experimentation is essential to establish where remanufacturing and repurposing make economic sense, helping reconcile environmental goals with the competitive pressures manufacturers face."

Of course, digital technologies run through every aspect of advanced manufacturing and its successful future, making this a core pillar. AI, digital twins and robotics can improve quality, support engineering, reduce downtime and help coordinate changing production requirements. Željko is quick to add that the roadmap places particular weight on trusted industrial data, interoperability and secure infrastructure, including accessible tools for smaller companies.

"It also looks towards humanoid robots and increasingly autonomous factories," he says. "These are emerging opportunities whose value must be established in industrial settings. Their inclusion reflects the need to prepare for technologies developing rapidly, while staying focused on their reliability and safety and the need for human oversight."

That human dimension forms the fifth pillar. With an ageing workforce and changing skills requirements, Europe must help people work effectively with more sophisticated equipment. Željko emphasises relieving physically demanding or repetitive tasks and allowing employees to contribute through supervision and other skilled activities. The roadmap couples technological adoption with continuous learning, workplace design and meaningful human control.

The five pillars reflect the range of industries that depend on manufacturing and the new markets European companies need to pursue. Željko argues that Europe cannot rely on its established strengths in its major sectors like automotive and aerospace alone. He sees opportunities in areas such as clean energy equipment, food processing, drones and humanoid robotics, where existing manufacturing expertise can help businesses develop new products and enter growing markets.

EFFRA's six working groups help translate these industrial needs into research and innovation priorities. Bringing together expertise from across its membership, they examine productive and flexible manufacturing, circularity, people and AI, energy, transport, and exploitation of research results. Their contributions inform the roadmap, connecting its wider ambitions with the technologies, skills and practical support companies need

Turning those priorities into wider industrial uptake is a particular focus of the working group on exploitation of research results. This requires planning from the outset for how a project's outcomes could be used, involving potential adopters throughout development and communicating results in ways that help manufacturers assess their relevance. For Željko, demonstration is essential: companies need to see how a technology performs under real production conditions and what advantages it offers before they will invest. Effective communication helps that evidence reach businesses beyond the project consortium, making exploitation an integral part of the project lifecycle.

That process depends on industrial partners willing and able to test new technologies in their own operations. Involving them early helps researchers understand production requirements and gives companies time to prepare the skills, resources and business case needed for adoption. "The reality check only comes if you have companies involved in projects who are also ready to bring the researchers to the factory floor," Željko says.

That is the argument for continuing a manufacturing partnership in the next framework programme – FP10. Željko highlights technology readiness levels five to seven, where technologies move through validation and demonstration towards greater industrial maturity. Bringing universities, research organisations and manufacturers together at this crucial stage allows practical requirements to inform development and gives businesses a clearer view of what adoption entails. "Getting these networks working properly will lead to far greater uptake of research outcomes," Željko adds.

"Collaborative programmes are key to this," he continues. "Public support creates an incentive for companies and specialists to work together, share costs and tackle challenges they might otherwise pursue separately. The next step is to strengthen the connection between those collaborations and wider industrial uptake."

The draft roadmap proposes complementary roles for the next framework programme, the European Competitiveness Fund and other instruments, connecting applied research with deployment, infrastructure and training. Its circular manufacturing proposals also call for monitoring at 12, 24 and 36 months after projects finish. "Following continued operation and market uptake would help establish whether demonstrators become lasting industrial capabilities, while giving subsequent projects evidence on which to build," argues Željko.

He believes SMEs need particular support to participate in these developments and benefit from the resulting technologies. "Many cannot sustain extensive research activity or devote large teams to evaluating and integrating unfamiliar technologies," he says. "So the roadmap envisages access to shared facilities, affordable tools and support that brings expertise within reach."

Željko also points to cascade funding, through which smaller organisations can receive support without joining a project's main consortium, and European Digital Innovation Hubs as ways of extending participation.

"Successful deployment depends on reaching companies beyond those receiving project funding," he says. "Demonstrations and results must be visible, understandable and relevant to potential adopters. EFFRA's Manufacturing Partnership Days provide one very important forum for that exchange. This gives businesses the chance to see project outcomes and explore technologies without first committing to a research consortium."

This wider network is an important part of Željko's case for EFFRA. Europe has strong regional industrial ecosystems and national research capabilities. A European partnership connects them, allowing knowledge and experience to travel between places that might otherwise remain relatively separate.



EFFRA's draft roadmap is driving the manufacturing agenda in Europe

"The work of EFFRA is certainly beyond the mere Made in Europe projects," he says. "Companies can use its activities for technology scouting; research organisations and national innovation bodies can compare approaches and establish relationships. Those connections create routes through which results can reach manufacturers outside their original project networks."

Preserving that capacity matters because, in Željko's assessment, Europe still possesses a broad industrial ecosystem: specialist machinery makers, production equipment suppliers, digital technology companies, researchers and skilled people. He points to businesses such as Bosch, Comau and Michelin as examples of the capabilities available. His concern is ensuring that Europe continues to master the technologies on which its production depends.

The roadmap therefore places that ambition in an international context. Other major economies actively support their industrial capabilities and emerging technologies. Europe also needs the capacity to develop and scale its own solutions, retain industrial know-how and build businesses around new opportunities. Its existing strengths provide a foundation, but dependence on external suppliers for critical materials, components and digital infrastructure leaves manufacturers exposed to disruption and decisions taken elsewhere. Meanwhile, an ageing workforce and shortages of engineering, digital and technical skills can constrain companies' ability to adopt new technologies and expand production.

Preserving Europe's manufacturing capability therefore requires continued investment in secure supply networks and European technology providers, alongside training, knowledge transfer and attracting new people into industrial careers.

Željko is clear that EFFRA is key to addressing these challenges. "Our role is to help companies draw on those European capabilities through research cooperation, demonstration and access to expertise."

But he draws a clear distinction between this practical role and legislative decisions on competition or ownership, which lie beyond EFFRA's remit. Its influence comes through creating opportunities for businesses to cooperate and grow. As he puts it, "through these programmes, we bring the experts together. Public support helps companies share the costs of developing and demonstrating technologies, building relationships that might not otherwise emerge and strengthening Europe's industrial capabilities in the process."

The case for EFFRA's next chapter rests on that practical connection. Its roadmap sets out a clear direction; its partnership and networks offer a means of bringing the necessary people and technologies together. "The measure of success will be how widely those capabilities are used to sustain competitive production and growing businesses across all of Europe," Željko concludes.

MATERIALS AND SUSTAINABLE PRODUCTION

MAGELLAN: Small magnets, strategic power

Permanent magnets may represent a market worth tens of billions, but they underpin industries worth trillions – and Europe imports 98 per cent of its rare-earth magnet demand, mainly from China. **Alexandre Damiens**, Programme Manager for permanent magnets and rare earths at Orano, explains how the 14-partners MAGELLAN project is demonstrating how better recovery, smarter motor design and European manufacturing could turn materials already in circulation into a source of industrial resilience.



Inert atmosphere chamber used for transferring magnets to different production stages

Permanent magnets are one of modern industry's most striking examples of a small component having a big impact. Found in products ranging from mobile phones and electric vehicles to wind turbines, they concentrate exceptional magnetic power into a relatively small amount of material. Their physical footprint may be modest, but their influence on the performance and even the viability of the finished product is enormous.

"The ratio between the magnetic energy and the volume or weight is unmatched," says Alexandre Damiens of Orano, the company coordinating the MAGELLAN project. "In an electric-vehicle traction motor, you have around two kilograms of high-grade permanent magnets, and you can move a car of two or three tonnes."

Permanent magnets create a persistent magnetic field without requiring a continuous electrical supply. In the high-performance neodymium-iron-boron magnets at the centre of the project, a precisely controlled composition and microstructure deliver exceptional magnetic-energy density. Motors and generators can therefore be made smaller, lighter and more efficient.

"The permanent-magnet market is worth tens of billions a year worldwide, but the markets that depend on it are worth much more – multiple trillions in fact," Damiens explains. "You have electric vehicles, drones, computers, offshore wind and more. It creates a lever which is more strategic than economic."

Their small footprint in a product's bill of materials can therefore conceal a strategic dependency. As Damiens puts it: "The weight, price or volume of a magnet does not represent its strategic value – and it is the strategic value of the magnet that is really valuable."

For Europe, that value is accompanied by an unusually concentrated supply risk. The European Commission estimates that 98 per cent of the EU's rare-earth magnet demand is met by imports, predominantly from China. Over several decades, China has developed a highly competitive ecosystem covering rare-earth processing, alloy production and magnet manufacturing. Europe does maintain extensive downstream capability in industries that consume magnets, but produces very few of the magnets itself. "It is a huge concentration," says Damiens. "China has developed an ecosystem, and it is very competitive too and Europe cannot match that."

MAGELLAN – MAGnets in rEsiLient suppLy chAiNs – was launched in 2024 as a Horizon Europe Innovation Action. Its 14 partners include industrial companies, research organisations, universities, recyclers, materials specialists, magnet manufacturers and automotive end-users.

"The project's goal is to kick-start a new European value chain specifically for permanent magnets," Damiens says. "As Europe is not endowed with large quantities of critical minerals,


Alexandre Damiens

recycling has to be part of the picture. Building that circular capability in Europe puts the project at the crossroads between the clean transition and the industrial stakes that we now call sovereignty."

For MAGELLAN, circularity is therefore the means by which sovereignty becomes a practical outcome. Europe may not control the primary supply of rare earths, but it can gain greater control over materials already circulating in its vehicles, turbines and electronic devices – provided they can be recovered and returned to high-performance uses. At present, neither the products nor the waste systems around them were designed to make that easy.

Lost by design

Much of the material installed during the rapid expansion of electric vehicles and modern wind power has not yet reached the end of its working life. "To recycle, you need to have end-of-life magnets," Damiens observes. "As the first wave of electric vehicles and wind turbines is not yet fully retired, that is a block for mass recycling."

Light e-mobility products with shorter service lives provide an earlier source, while larger volumes from electric vehicles will emerge over time. However, the availability of feedstock is only part of the problem.

A vehicle enters an established end-of-life system and contains a relatively concentrated magnet mass. Phones, headphones and other consumer devices are dispersed among millions of users and contain only a few grams of permanent magnets. Unless those products are collected, dismantled and sorted specifically for their magnets, the material disappears into a much larger waste stream.

Once that happens, their functional and strategic value is effectively diluted away. Europe may have imported the rare earths and paid for the energy, processing and expertise required to make them into high-performance magnets, but the materials do not return to an equivalent use.

Product design and construction create a further barrier, as magnets in electric motors and consumer devices are

frequently glued into place, so recovering them means removing the drive, motor, rotor and individual magnets, with each stage adding time, cost and a risk of contamination. "You have the time effect, the organisation of the waste streams and the way the magnets are incorporated into the device," Damiens explains. "At the moment, they are not easy to dismantle".

"And even after recovery, the magnets are not interchangeable," he continues. "Waste products may contain ferrite, aluminium-nickel-cobalt or neodymium-iron-boron magnets. Different grades of neodymium-iron-boron material can contain different proportions of elements, including valuable heavy rare earths such as dysprosium and terbium. Coatings, adhesives, oxidation and carbon contamination can also affect the recycling route."

"You do not have only one steel, you have many types," Damiens adds. "It is the same with magnets. You need to sort the different types and grades."

Recycling facilities need sufficient feedstock and assured customers, yet buyers hesitate before capacity, quality and price have been demonstrated. MAGELLAN must therefore address the entire system rather than a single recycling process.

Closing the loop

MAGELLAN's response to this is organised around three connected pillars: make, use and recover. "In making, we try to use less material and introduce partial substitution," says Damiens. "In use, we reduce the amount needed per kilowatt and facilitate future dismantling. Recover is about better recovery, reuse and recycling at the end. It is a virtuous circle."

"The make pillar of the project explores more efficient use of rare earths, the incorporation of secondary material and the partial replacement of the most critical elements," he continues. "Use pillar focuses on electric-vehicle traction motors that require less magnet mass for the same output and can be dismantled more easily. Recover covers collection, extraction, sorting, reuse and recycling."

"Some partners are more cost-oriented and others are more focused on high performance," Damiens adds. "Our task is to get these worlds talking to each other."

Within the project, shared milestones bring these areas of expertise together. Developing magnet specifications connects materials and motor specialists, while designing extraction processes involves dismantlers, recyclers and rotor engineers.

"Everyone has to reach the milestone, and we need each other," says Damiens. "The recovery process begins with end-of-life vehicles and light e-mobility products sourced in Europe. After the motor or drive has been removed, the rotor is extracted and the magnets separated. Characterisation then determines the most suitable destination for each item."

One sorting system may use laser spectroscopy to analyse composition. Its response can reveal factors such as oxidation, carbon contamination and the presence of valuable heavy rare earths. Artificial intelligence helps to refine the relationship between the spectrometry results and the robotic sorting decision. "The magnets move on a conveyor, the system analyses the response and the robot sends each one to the right place," Damiens explains. "AI is being introduced in that type of operation to make the targets and the spectrometry response more accurate."

For suitable neodymium-iron-boron material, MAGELLAN favours a short-loop route that keeps the recovered material in metallic form rather than returning it to separated rare-earth oxides, so-called long-loop route. "You do not need to go back to the oxide state and through the entire value chain," says Damiens. "You avoid energy-intensive and chemically intensive steps. That is why it is called a short loop."

Recovered scrap is remelted and strip-cast into ribbons or flakes. Hydrogen decrepitation creates a coarse powder, jet milling refines it, and pressing and sintering produce dense magnets. Composition, particle size, temperature and microstructure must be controlled throughout.

"These are industrial and metallurgical steps that need to be mastered in Europe to bridge the skills gap," Damiens says. "That is also embedded in the project."

Powder Injection Moulding offers another way to improve material efficiency. Conventional production may begin with a rectangular block that must be machined into the required geometry, creating cutting waste. Near-net-shape manufacturing forms a magnet much closer to its final dimensions before sintering.

"You can shape the trapezoid directly instead of starting with a rectangle and cutting it," Damiens explains. "It avoids the cutting process at the end and avoids creating another waste stream."

"Shape also influences the magnetic flux within the motor," he continues. "A carefully designed magnet can therefore reduce manufacturing waste while helping the motor produce more output from each kilogram of material. "Meanwhile, composition involves its own trade-offs. Cerium and lanthanum can replace part of the more critical neodymium, but they offer weaker magnetic properties."

MAGELLAN is targeting high-performance magnets containing 25 per cent recycled feedstock rather than pursuing recycled content at the expense of function. "It is very hard to maintain a high rate of secondary material and stay at the state of the art in performance," Damiens says. "We do not want to recover an electric-vehicle magnet and downcycle it into a low-value use. That is not meaningful for an autonomy strategy."


Final sintered magnet blocks

Motor designers, meanwhile, approach the challenge from the other end. Their concern is not the chemistry of the magnet in isolation, but whether it supplies the required flux and performance. "The automotive partners do not think in terms of the magnet's chemistry; they think in magnetic properties," says Damiens. "They want the flux, and their aim is to use less magnet for the same power."

Reversible assembly and the elimination of permanent glueing can then make those magnets easier to access at end of life. Design for performance and design for recovery become the same challenge viewed at different points in time.

Performance retained

MAGELLAN is working towards a TRL7 demonstration in which recycled and near-net-shape magnets are integrated into electric-vehicle traction motors. These will be tested against reference motors.

The assessment covers magnetic and motor performance, magnet use per kilowatt, assembly and dismantling, alongside life-cycle and cost analysis. Although research-scale costs cannot reproduce the economies of an established industrial plant, the demonstration should show where costs arise and how performance changes across the process.

The 25 per cent target is intended to preserve the magnet's original function: "We do not want a downcycling approach," Damiens stresses. "We set a trade-off between the incorporation of secondary materials and high performance." "This is recycling measured by retained function rather than simply by weight diverted from waste. The rare earths remain in a use that continues to exploit the properties for which they were originally mined, refined and processed."

Shortening the process also brings environmental benefits. In some scenarios, the project's life-cycle work indicates that using 25 per cent secondary material could halve process-related carbon emissions compared with a conventional production route based in China. "Because it is a short loop and you are not going far back through the process, with 25 per cent secondary material you can have up to a 50 per cent reduction in CO2 emissions," says Damiens.

Results will vary with feedstock, energy mix and manufacturing scale, but the finding shows why recycled content and process design must be assessed together.

Alongside the prototype, MAGELLAN is reconnecting industrial skills and developing guidance for motors and rotors that allow magnets to be removed more easily. The project also plans to publish a white paper around its final conference in Brussels in February 2027.

The intention is to leave behind specifications, process knowledge and working relationships that manufacturers and recyclers can develop after the funded programme ends. Together, these outcomes could reduce Europe's demand for primary extraction abroad and put material already within its economy back to work. Damiens is careful, however, not to overstate what one project can achieve. "It is a step, not the full way," he says, "but it is a

step for industrial sovereignty on a component that today is almost entirely imported."

From pilot to production

MAGELLAN is working with physical products and real dismantling operations to understand what implementation would involve outside the laboratory. "It is a real test; it is not a lab thing," Damiens says. "We dismantle motors, so we have the cost of access to the end-of-life magnet in mind."

"Removing the drive, rotor and magnets takes time, and every minute affects the economics," he continues. "A technically elegant recycling process cannot succeed if the material arrives at its gate carrying an uneconomic recovery cost."

Beyond dismantling lies the transition from demonstration capacity to bankable demand. "We need offtake, but that is a chicken-and-egg problem," Damiens explains. "If you do not have capacity, there is no offtake; if you have no offtake, you do not build capacity."

Once European magnets have been demonstrated and qualified, automotive manufacturers and Tier 1 suppliers must create durable demand. Because magnets are developed for particular motors, offtake involves technical qualification, price, volume and contract duration rather than a simple spot purchase. Damiens sees scope for buyers to divide their sourcing between a business-as-usual majority and a strategic European share.

"An OEM might keep 80 per cent of its normal supply and take a balanced approach on 20 per cent to de-risk a sensitive component," he says. "Offtake is not only price. There is price, length and volume, and those parameters take time to align.

The principle here is that resilience can be improved without attempting to separate completely from the existing global supply chain." Standards could lower the cost of circularity by allowing recyclers to identify a magnet's composition, grade and condition without reconstructing its history through testing. "Labelling helps you have a more efficient recovery process," Damiens says. "With a QR code or a new standard approach, you have better knowledge of the secondary feedstock."

MAGELLAN is contributing to standardisation discussions at European and international levels, while sorting remains essential for the unlabelled legacy products that will dominate the waste stream in the near term.

On waste classification, MAGELLAN and eight other Horizon Europe projects representing more than 100 entities are calling for dedicated rare-earth and permanent-magnet codes within the EU list of wastes. "Waste codes are the way to track waste streams," Damiens explains. "We tried to shape the discussion around what industry needs."

Such codes would make the stream visible and traceable, supporting collection, statistics and investment. European policy is moving in the same direction. The Critical Raw Materials Act addresses recycling and diversification, while the Net-Zero Industry Act, RESourceEU, proposed Industrial Accelerator Act and planned Circular Economy Act seek to stimulate demand for European-made, low-carbon and secondary materials.

For Damiens, however, policy direction is only the beginning: "Things are moving, and there are good signals," he says, "but signals are not sufficient. Those signals must ultimately become investable demand, usable standards, accessible waste streams and contracts capable of supporting industrial capacity."



Strip-casting of permanent magnet alloy

A system, not a component

MAGELLAN's clearest lesson is that circularity cannot be delegated to recyclers at the end of a product's life. Decisions about bonding, labelling, material quality and future demand determine whether a magnet can return to high-value use.

These choices shape the viability of the wider system. Sovereignty depends less on where a finished product is assembled than on whether the surrounding industrial system can respond when supply is disrupted.

Europe therefore needs enough capability in collection, separation, alloy preparation, powder processing, magnet production, motor design and testing to create credible alternatives. "There are a lot of different competencies that we are lacking at the moment that we are trying to gather together through MAGELLAN," says Damiens. "That's why the project has 14 partners – because you have the engine designer, the people focused on pressing the magnet and others focused on extracting it."

That value is easily missed when magnets are treated as small bill-of-materials items. Europe may struggle to match prices created by decades of scale in China, but buyers also gain qualified alternative supply, shorter processing routes, retained expertise and reusable material.

"If you want a magnet at the Chinese price in Europe, that can be very difficult," Damiens acknowledges. "But if you look at it as diversification, localised content and a de-risking approach, it is different."

A de-risked future

The project's vision of the future is more pragmatic than complete self-sufficiency. China's expertise, scale and competitiveness will remain important to the global magnet market and so Europe's objective is to create enough alternative capability to absorb shocks, qualify secondary material and prevent a shortage of one small component from stopping much larger industries. "We need to find the right amount of industrial capabilities and capacities that can buffer excess pressure," Damiens says. "It will be somewhere in between, and that is why I speak about de-risking more than decoupling."

The remaining milestones will show whether the loop closes: recovered magnets must return to high-performance use, near-net-shape production must cut losses while meeting motor requirements, and rotor designs must use less material while releasing it more easily at end of life. The resulting cost and life-cycle evidence must be credible to buyers.

MAGELLAN's Brussels conference and white paper will help translate the results into industrial and design guidance, but the partnerships will need to continue beyond the project. Success would leave Europe with a pool of rare earths already paid for through imports, rebuilt manufacturing knowledge and less exposure to a single dominant source.

It will also have created an additional option for automotive, energy, digital, aerospace and defence industries whose combined value is measured in trillions. "In a de-risking approach, it is not only the price that matters," Damiens concludes. "It is also the freedom of action you preserve."

PROJECT INFORMATION

Project Title
MAGELLAN – MAGnets in rEsiLient suppLy chAiNs

Project Objective
MAGELLAN aims to engineer and demonstrate at TRL7, a sustainable and cost-effective rare-earth-element (REE) high-performance permanent magnets (PM) value chain solution for electric vehicle traction motors (EVTM)

Project Duration and Timing
39 months (01/01/2024 – 31/03/2027)

Project Funding
Horizon Europe, Total budget: €7 578 190,26, EU contribution: €6 503 323,63

Project Partners
Orano, Magneti Ljubljana, Centro Ricerche Fiat, Comet Traitements, CEA, IFP Energies nouvelles, Universite de Liege, Fundacion Tecnalia Research & Innovation, Technische Universitat Darmstadt, Geologian Tutkimuskeskus, Lomartov, RTD Talos, EIT RawMaterials, BLC – The Battery Lifecycle Company.

MAIN CONTACT

Name
Yiannis Paphitis



Phone
+357 22454333

Email
info@magellan-horizon.eu

Web address
https://magellan-horizon.eu

This project has received funding from the European Union's Horizon Europe research and innovation program under the grant agreement **No 101138260**.

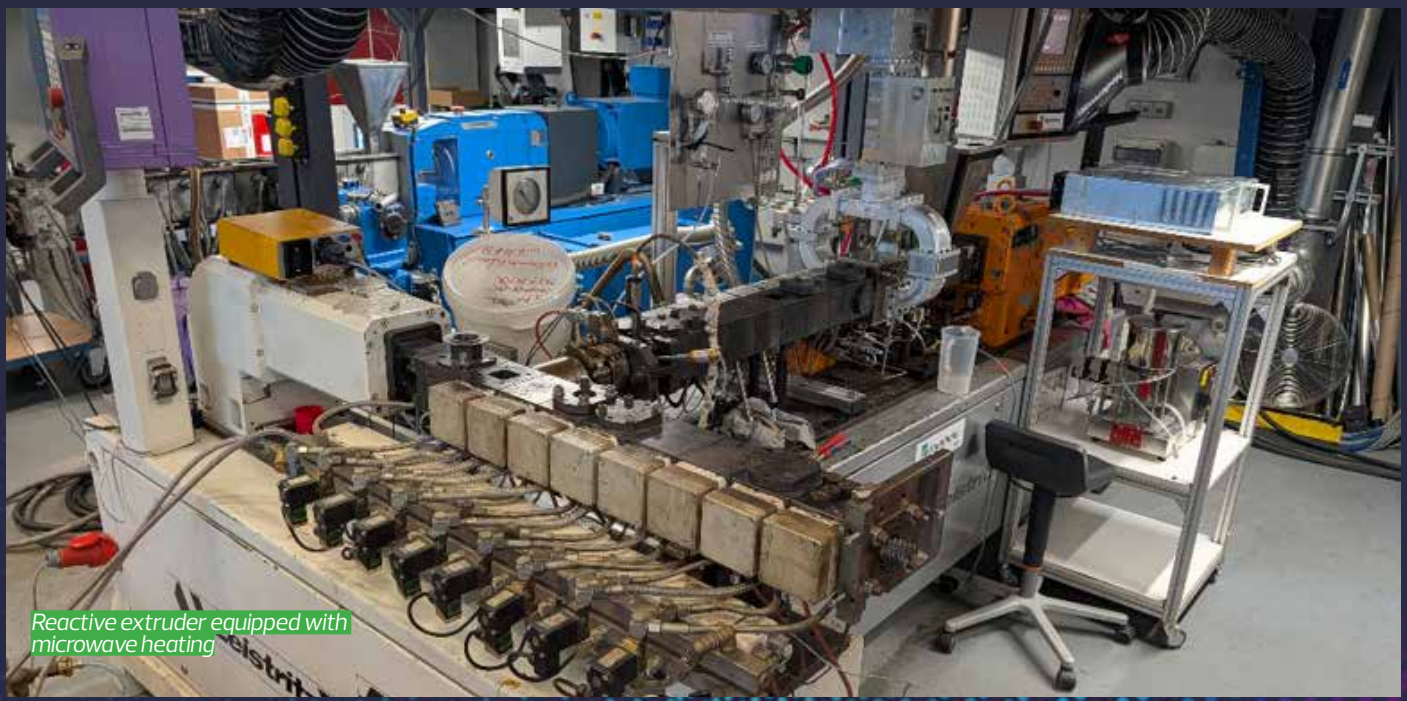
MATERIALS AND SUSTAINABLE PRODUCTION

SIMPLI-DEMO: From batch certainty to continuous opportunity

Many speciality-chemical and pharmaceutical processes still rely on batch production because solid-laden and highly viscous materials remain difficult to handle continuously. SIMPLI-DEMO is developing and demonstrating modular continuous systems that combine advanced reactors with ultrasound and microwave activation, real-time process analysis, dynamic modelling and automated control. Professor **Tom Van Gerven** of KU Leuven and **Dr Evangelos Delikonstantis** of AristEng Sarl explain how the project aims to make European chemical manufacturing more flexible, efficient, electrified and competitive.

For **Tom Van Gerven**, Professor of Chemical Engineering at KU Leuven and project coordinator for SIMPLI-DEMO, the motivation for the project's work was relatively simple. Chemistry and biology have advanced enormously, but the process technologies used to turn those advances into manufactured products have not always kept pace. "We've improved so much over decades in terms of chemistry, biology and the fundamental sciences," he says. "But process technology knowledge, process technology investment and novelty are all lagging behind. All should be following the improvements being made in chemistry and biology."

Many speciality-chemical and pharmaceutical industries still rely on batch production to process the chemicals they need. A vessel is filled with reactants, the material is mixed, heated or otherwise processed, the completed product is removed



Reactive extruder equipped with microwave heating

and the cycle begins again. Batch equipment is familiar, comparatively flexible and able to accommodate different products, formulations and production schedules. It can also handle the solid-laden and highly viscous mixtures that occur frequently in these sectors.

It has its limitations, however, and these arise from the same stop-start structure of the process, with the equipment having to be repeatedly filled, heated, cooled, emptied and prepared for another cycle. Production conditions can vary between batches, while process information often becomes available only after laboratory samples have been analysed. Energy is lost during heating, cooling and periods of inactivity, and the equipment can occupy considerable space relative to the amount being produced.

Continuous processing offers another model, with the reactants entering a reactor while product is removed continuously, allowing the process to remain under stable operating conditions. Continuous technologies have consequently become established for many homogeneous systems involving liquids or gases.

However, processes containing suspended solids or very viscous materials are considerably more difficult to operate continuously. Solids can accumulate, deposit and block tubular reactors, while thick or paste-like materials are difficult to convey and mix. Continuous systems can also appear less adaptable than batch plants when a manufacturer needs to change operating conditions or produce several product grades using the same equipment.

"The types of processes that are lagging in this transition are particularly those which contain viscous streams or solids," Van Gerven explains. "Certainly in the speciality and pharmaceutical industries, these are very important types of processes."

SIMPLI-DEMO was conceived to address both limitations: the physical difficulty of processing challenging materials and the perceived lack of flexibility in continuous production.

"Another key motivation of the project is to support the chemical industry in the transition from batch to continuous," adds **Van Gerven**. "We believe that is important because it will provide a competitive edge for European industry."

The objective is therefore broader than replacing one batch vessel with one continuous reactor. SIMPLI-DEMO is developing modular production systems in which continuous reactors are combined with ultrasound or microwave energy, online process analysis, dynamic modelling and automated control. Alternative energy can be delivered at selected points in the process and adjusted rapidly, giving operators additional ways to control mixing, mass transfer, heating and product characteristics.

The approach is being evaluated through four industrial cases selected to represent the materials that have proved most resistant to continuous processing. Two concern solids: crystallisation of zeolite materials by Arkema and the crystallisation of precursors for pharmaceutical ingredients by GE Healthcare. Two involve highly viscous polymerisation processes at Coatex and Megara Resins, including the production of non-isocyanate-based polyurethanes.

Together, the four cases test whether continuous, modular and automated production can be extended beyond ideal liquid-phase processes to the difficult chemistries that industry still commonly operates in batches.

Making continuous flow work

SIMPLI-DEMO combines two process-intensification platforms suited to these demanding material streams. A continuous oscillatory baffled reactor, or COBR, contains internal baffles and uses a pulsing flow to create effective mixing even when the net movement of material through the reactor is slow. This allows longer residence times to be achieved within a relatively compact system and makes the technology particularly relevant to solid-laden processes. A reactive extruder, or REX, uses a screw or twin screw to convey, mix, knead and react highly viscous or paste-like compounds that would be difficult to move through a conventional tubular reactor. These reactors can then be combined with ultrasound or microwave activation.



Sonotrode integrated in an extruder barrel

"Ultrasound is mostly used for better mixing and better mass transfer, whereas microwaves are mostly used for better heating," says **Van Gerven**. "Mass transfer, heat transfer and mixing are the three fundamental transport phenomena. With ultrasound and microwaves, we are able to intensify and improve all three of them."

Because both forms of energy can be applied selectively and adjusted rapidly, they provide additional control over the process. Both can be adjusted or stopped much more quickly than conventional heating systems. These new control levers also introduce additional engineering complexity: the ultrasound and microwave modules must be physically integrated with the reactor, while their power, timing and interaction with changing process conditions must be coordinated through the automation and control system.

"Obviously, you increase complexity," **Van Gerven** acknowledges. "You bring in additional equipment that has to be integrated not only from a hardware perspective, but also from the perspective of process control and automation."

Electrification nevertheless creates a form of responsiveness that conventional steam heating cannot. A steam jacket retains heat after the supply is turned off, while microwave energy can be stopped almost immediately. **Van Gerven** sees a longer-term opportunity in that low inertia: "A plant could increase the use of ultrasound and microwaves when electricity is inexpensive, reduce it when prices rise and potentially alter the grade being produced," he says.

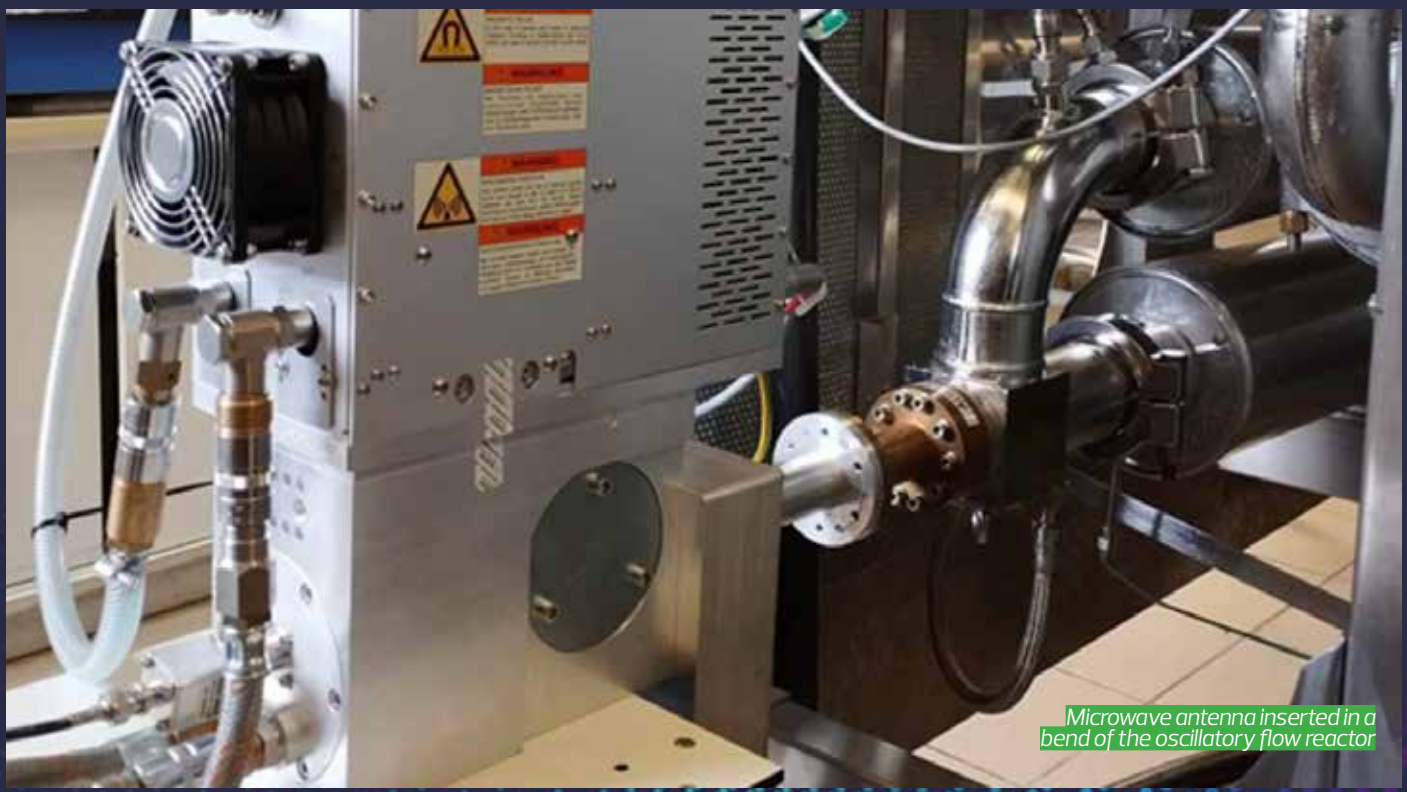
"Flexibility would then mean more than switching recipes; it could include responding dynamically to energy availability and cost."

Turning time into space

SIMPLI-DEMO's development sequence began with batch experiments to understand the chemistry in enough detail to design the continuous system. The team clarified each reaction, developed modular COBR and REX configurations, integrated ultrasound or microwaves, and then brought in process analytics, automation and control. Two use cases inherited knowledge from the earlier SIMPLIFY project; two additional cases tested how quickly that experience could be transferred.

The bridge from batch to continuous is best captured by **Van Gerven**'s phrase "moving from time to space". In a batch vessel, an action might occur two or four minutes after the process starts. In a tubular reactor, that sequence becomes a location. "If batch results show that it is better to use ultrasound at the end of a crystallisation process, for example to deagglomerate the crystals, then in a continuous system you should use it at the end of the reactor," he explains. "The end of your reactor reflects the end of the process. You are moving from time to space. That is how we optimised the location, power and timing of ultrasound and microwaves in batch and translated that to continuous.

Once the action has a location, the process still needs rapid feedback. Offline sampling is too slow when material



Microwave antenna inserted in a bend of the oscillatory flow reactor

continues to flow while a laboratory result is pending. **Van Gerven** puts the practical problem plainly: ''By the time you have the results, so much product has been wasted, so you need very quick online sensors.''

A near-infrared process analytical technology was therefore integrated in the Coatex case to follow yield or product conversion online. A common data layer allows the different levels of automation to communicate, while dynamic models connect operating parameters to efficiency and product-quality indicators.''

Evangelos Delikonstantis, Managing Director at project partner AristEng Sarl, describes the difficulty here: ''Microwave and ultrasound activation create local, transient and non-uniform conditions that are extremely demanding to represent using physical models alone. We therefore explored machine-learning approaches trained on experimental data to identify trends and predict whether changing a parameter will improve the process.

''Those models can serve as a sort of digital twin, linking critical parameters with important KPIs,'' he continues. ''Based on experimental data and customised experiments, they make predictions or correlations between critical operating parameters and KPIs such as efficiency and product quality.''

Another layer of complexity comes at the plant level because a percentage gain in reactor yield is not automatically valuable if it adds disproportionate cost or environmental burden elsewhere. The project therefore linked experimental results to economic and environmental assessments, helping the consortium identify which variables could materially improve the performance of the process as a whole.

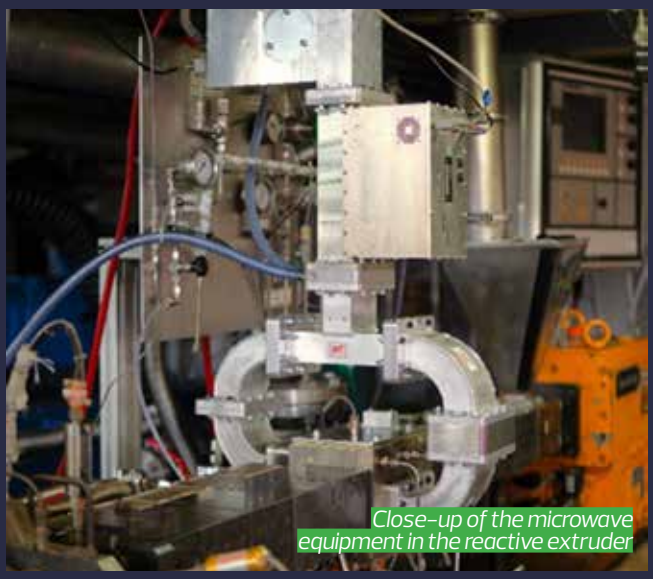
''It's a multi-layered approach,'' **Delikonstantis** adds. ''The equipment, its live behaviour and its wider industrial consequences are assessed together.''

Industrial pilot scale

As SIMPLI-DEMO approaches completion, all four cases have either been tested in a relevant industrial environment or are close to that milestone. The project's goal is industrial pilot-scale demonstration, not yet full production-scale deployment.

The Arkema crystallisation case provides a particularly clear illustration of what can change even when the underlying chemistry remains the same. Moving to a COBR integrates an upstream ageing step into the continuous operation, reducing the number of separate process stages. It also makes recovery of the crystallisation mother liquor more practical.

In batch operation, that liquid had to be collected in large quantities and treated separately, disrupting the process and making disposal more economical than recovery.



Close-up of the microwave equipment in the reactive extruder

Continuous flow changes that balance. **Delikonstantis** reports approximately 40% lower consumption of the auxiliary materials associated with the mother liquor, rather than a reduction in the critical raw materials used to make the catalyst. He also reports energy savings of around 20–25%, principally through avoiding repeated heating, cooling and downtime.

At the reactor level, the process has been electrified, removing the need for fossil-fuel-generated steam in the reaction zone. Together, those changes produced an estimated reduction of up to 10% in the case's overall carbon footprint.

''20-25% energy savings, 40% lower materials as inputs, and no natural-gas combustion – only green electricity,'' he summarises. ''Overall, we've seen that we can reach up to a 10% lower carbon footprint.''

The result also exposes a strategically useful limit in that if the principal raw materials dominate a product's carbon footprint, then improved energy efficiency and reduced auxiliary consumption can only go so far without changing the chemistry itself. ''Process intensification is powerful, but lifecycle analysis prevents it from being presented as magic,'' says **Van Gerven**.

Other findings have opened fresh technical directions. Work by the team at the National Technical University of Athens showed that cycling microwave temperature, rather than holding one constant setpoint, could deliver productivity increases of roughly 20–40% in some crystallisation and polymerisation experiments.

Ultrasound's benefit in one polymerisation case appeared mainly as local heating, contrary to the expectation that mixing would dominate. For crystallisation scale-up, the relevant quantity proved to be total ultrasound energy –

power multiplied by contact or irradiation time — rather than power alone.

"This might look very logical," **Van Gerven** notes, "but if you look at the literature, it is often not compared. It is mostly compared at the level of ultrasound power, not ultrasound energy.

"And that is the kind of modest insight that can determine whether a laboratory effect survives at larger scale."

The SIMPLI-DEMO platform

Across the four cases, the project has demonstrated that COBR and REX platforms can be adapted to solid-laden and highly viscous mixtures. That does not make them plug-and-play. Each chemistry still requires its own configuration, operating window and control strategy. The transferable asset is a set of reactor platforms, energy modules, modelling methods and integration knowledge from which a tailored process can be built.

"Fine-tuning is always necessary for the respective chemistry and reactions," **Delikonstantis** says. "But these two process-intensification technologies can be applied to highly viscous and solid-based mixtures, whether it is a food additive, an active pharmaceutical ingredient or a polymer."

Partners are already considering food and feed, environmental processes and CO_2 capture. Wider adoption, however, involves people and institutions as well as hardware. Engineers and operators need training in equipment that many chemical-engineering graduates have never encountered. Validation, regulation and site requirements will remain process-specific.

SIMPLI-DEMO has used company questionnaires and an industrial stakeholder panel to map these issues qualitatively and expose the questions future adopters will have to answer. The project does not prescribe a single route to market: commercialisation and intellectual-property decisions will remain with the individual partners.

Delikonstantis is candid about the skills gap. "I was never exposed to such high-tech or novel pieces of chemical equipment during our studies or training, and I believe most of my peers and colleagues were not exposed either," he says. "Customised training for workers or engineering teams might be needed.

"Demonstration in the future must, therefore, be accompanied by education," he adds. "The sector must understand not only the advantages of this technology, but where it is genuinely applicable and how it should be configured to deliver worthwhile results."

The industrial reality gap

The project's most important lesson may be that an established industrial process is not necessarily a fully understood process. Companies know enough chemistry to manufacture a compliant product under existing conditions. Changing the reactor introduces operating conditions that may never have been mapped in the literature or inside the company.

"You start with less knowledge than you would hope for," says **Van Gerven**. "When you go to another type of technology, you move into other operational conditions which have not always been mapped up front. We took more time than anticipated to understand the chemistry, find the optimal product specifications and operating conditions before we could go to the technology development.

"That 'industrial reality gap' slowed the progress we were able to make towards some of our initial ambitions," he continues. "And yet spending the time was not a failure; it was what made the demonstration credible." A second



Pilot installation of the continuous oscillatory baffled reactor

lesson came from choosing cases that already achieved high yields in batch processing. Such cases were industrially relevant, but left less room for dramatic yield improvements. **Delikonstantis** suggests that future demonstrations might also target chemistries limited by equilibrium, where local microwave or ultrasound effects could reveal more of the technology's potential.

Van Gerven recalls GE Healthcare's conclusion: "Maybe the improvements for this particular case-study are not spectacular, but we now know what to look for within our company."

That knowledge is difficult to express as one headline KPI, but, he argues, "it is as important as those tangible results."

The future

Much of the new infrastructure sits inside individual companies, which limits the ability for other manufacturers to test their processes before committing capital. **Van Gerven** advocates open-access pilot facilities where a company could evaluate COBR or REX, with or without ultrasound and microwaves, together with the necessary analytics, automation and control.

Facilities already exist for aspects of the batch-to-continuous transition, **Van Gerven** observes, but alternative-energy activation is "further from the market" and has rarely been central to those installations. KU Leuven is currently building such infrastructure.

He believes the project results and such increased access to testing could have a profound impact on European industry. "All these technology-improvement projects should help keep the chemical industry in Europe," he says, pointing to the infrastructure, suppliers and knowledge being developed across companies and universities.

Delikonstantis goes further: "It is technology made in Europe, and most importantly the know-how development is here. Instead of importing technologies, it is better to export know-how. Europe should capitalise on all this knowledge and process-intensification technology developed over the past years."

Neither argues that intensified continuous reactors should replace conventional equipment everywhere. The more convincing vision is selective: identify the chemistries for which compact continuous reactors, precisely delivered energy and live predictive control offer a decisive advantage.

"I would be very happy to see those technologies not replace state-of-the-art or conventional technologies, but find the position they deserve," **Delikonstantis** concludes. "I'm sure there are applications where they can perform much better. We are trying to find and investigate those applications. That is the excitement I feel every time I deal with process-intensification technologies."

Van Gerven believes this may be SIMPLI-DEMO's most important outcome: not one universal solution, but a growing set of difficult industrial processes for which continuous, electrified and intelligent production becomes the practical choice.

PROJECT INFORMATION

Project Title
Demonstration of Sonication and Microwave Processing of Essential Chemicals (SIMPLI-DEMO)

Project Objective
SIMPLI-DEMO aims at strengthening the chemical process industry – in particular the specialty chemicals and pharmaceuticals industries – to produce materials and chemicals in a sustainable and competitive way by moving from batch to continuous and modular production with flexibility being ensured by the application of alternative energy forms (microwaves and ultrasound).

Project Duration and Timing
1 October 2022 – 30 September 2026

Project Funding
Horizon Europe, 9.4 million euro funding

Project Partners
Ku Leuven (Belgium)
Technical University Of Dortmund (Germany)
National Technical University Of Athens (Greece)
Fraunhofer Institute Of Chemical Technology (Germany)
Heriot-Watt University (United Kingdom)
Arkema France Sa (France)
Ge Healthcare As (Norway)
Coatex Sas (France)
Megara Resins (Greece)
Weber Ultrasonics Ag (Germany)
Muegge Gmbh (Germany)
Aristeng S.a.r.l. (Luxembourg)
Leikon Gmbh (Germany)

MAIN CONTACT



Name
Tom Van Gerven

Email
Tom.vangerven@kuleuven.be

Web address
https://simpli-demo.eu/

SIMPLI-DEMO
DEMONSTRATION OF SONICATION & MICROWAVE PROCESSING OF ESSENTIAL CHEMICALS

This project has received funding from the European Union's Horizon Europe research and innovation program under the grant agreement **No 101058279**.

INSIGHT FEATURE

Why "Made in Europe" must include the UK

As a British company dedicated to promoting European research and manufacturing success, we believe excluding the UK would weaken the industrial base this policy is intended to strengthen.

By **William Davis**, Managing Director, Insight Media

Europe needs more manufacturing capacity, stronger supply chains and the confidence to invest in its industrial future. Creating fresh barriers between British and EU manufacturers would work against all three aims.

That is the central absurdity threatening the EU's "Made in Europe" agenda. A policy intended to strengthen European industry could end up weakening an industrial base that already spans the English Channel.

The European Commission's proposed Industrial Accelerator Act would use public procurement and public support to encourage low-carbon production and European manufacturing. The ambition is understandable and one we support. Governments have legitimate reasons to protect strategically important capabilities and ensure that public investment creates lasting industrial value.

But the definition of "eligible production" matters enormously here. If British factories and components are disadvantaged in sectors where UK and EU businesses operate as an interconnected system, European preference when "European" is defined only as the EU, quickly becomes European self-harm.

Nowhere is this clearer than in the car industry. Analysis commissioned by the Society of Motor Manufacturers and Traders estimates that UK automotive manufacturing supports €24 billion of economic activity and 250,000 jobs across the EU. Those are European livelihoods directly connected to British production.

A policy that makes a British-built vehicle less competitive can therefore damage suppliers inside the Union. It can reduce demand for their components, weaken production volumes and complicate investment decisions. There are underlying industrial relationships behind these facts and redrawing the eligibility map won't change that.

When Britain left the European Union, it did not stop being part of Europe's manufacturing economy.



There has been encouraging movement towards a solution by both sides. Reuters reported on the 27th of October that business minister Jonathan Reynolds believed Britain and the EU were "finding a way through" the concerns. That is welcome. However, manufacturers need an agreement they can use when making investment decisions, rather than reassurance that negotiations are progressing.

Legal considerations

The proposals do not amount to a blanket exclusion of British goods, and provisions already recognise UK content in some circumstances. The problem is that sector-specific restrictions could undermine that recognition precisely where industrial relationships are deeply integrated.

Britain should press for clear equivalent treatment for genuinely UK-manufactured products and components in the relevant procurement and support rules. UK manufacturers should receive consistent treatment under the rules, without having to rely on special exceptions.

Of course, that requires practical commitments from London too. Reciprocal access should be part of any settlement. British ministers cannot credibly demand that EU schemes recognise UK production while treating European partners as outsiders in their own programmes.

Concerns about Britain circumventing any new rules should also be addressed head-on. Goods should not qualify as British merely because they pass through a UK port or acquire a convenient label. Clear rules on where goods are made, checks on actual production and joint enforcement will ensure that only genuine manufacturers benefit.

Any wider conditions involving trade policy towards China must be negotiated openly and assessed against the consequences for industry on both sides. Neither slogans about sovereignty nor vague promises of partnership will resolve those choices.

Manufacturers on both sides of the Channel should make their case together. They should show which supply chains, investments and partnerships would be disrupted, and how excluding Britain would also hurt businesses within the EU.

Europe's industrial strength depends on its ability to turn shared expertise into products that can compete globally. That demands scale, investment and cooperation across a trusted industrial network.

A well-designed policy can defend that network against unfair competition. A poorly designed one will divide it. Made in Europe should strengthen the factories Europe already has – including those in Britain.



When Britain left the European Union, it did not stop being part of Europe's manufacturing economy.

SMARTER, MORE FLEXIBLE FACTORIES

REED: Turning Europe's large-part manufacturing capacity into a service

Large components can take weeks to manufacture. A company may own a machine large enough for the work, but that does not necessarily make it a qualified supplier: it also needs the specialist knowledge to quote for the job and produce the component to the required standard. The €5.3 million Horizon Europe REED project is developing a Manufacturing-as-a-Service (MaaS) platform that connects such manufacturers with digital twins, monitoring tools, equipment and specialist services while keeping valuable production know-how under its owner's control. Project coordinator **Pedro Maria de la Peña Tejada**, **Alex Iglesias** of Ideko, and **Vasiliki-Eleni Provopoulou**, **Panagiotis Kourempanas** and **Konstantinos Chisiridis** of Netcompany, explain how REED aims to turn Europe's concentrated large-part expertise into a more responsive value network.

Manufacturing-as-a-Service (MaaS) platforms give companies digital access to distributed manufacturing capacity, equipment and services. Customers can define their requirements, identify suitable suppliers, compare quotations and coordinate production through a shared digital environment. Manufacturers can offer specialist assets and expertise beyond their established customer relationships. This model is most effective where requirements and capabilities can be described in standardised terms and matched to an established process.

The REED project is addressing a harder proposition: whether MaaS can work for large, complex parts that are expensive, produced in very small numbers and dependent on highly specialised manufacturing knowledge.

The difficulty becomes apparent when a customer requests a quotation. For a standard component, the requirement can often be matched with a supplier's equipment and a production route it already understands. For a large or one-off part, the equipment requirement is only the first test. Before responding, the supplier must determine how the workpiece will be handled, supported and aligned; which machining strategies will achieve the required tolerances; and whether it can commit to the requested quality, cost and delivery time. REED is seeking to bring these pre-production considerations into the MaaS model, rather than treating quotation as a simple match between a drawing and an available machine.

"Some of our partners take two or three weeks to produce one piece," explains REED project coordinator Pedro Maria de la





Peña Tejada, Chief Innovation and Emerging Technology Officer at Ayesa. "They are very long, very complex processes because they have a lot of operations. Production volumes may amount to only two, three or four units a year, while the size and value of each component mean that even a modest improvement in quality can have a substantial financial effect.

"Maybe they have the machine, but they don't have the expertise to produce it," de la Peña continues. "The idea is to minimise the risk: to help them make the quotation and, when they have the order, to produce on time without any problem with the quality."

REED is therefore developing a platform through which customers can request parts, manufacturers can describe their capabilities, and simulation, production intelligence and specialist tools can convert an unfamiliar requirement into a credible production plan. The project addresses the gap between owning suitable equipment and possessing the complete capability needed to accept and deliver the work.

The project is testing this proposition in three distinct industrial settings. **Komatsu** represents the yellow-goods sector with an excavator arm; **Maschinenfabrik Wüstwillenroth** (WWR) is working on an upper component for a large press; and **Goimek** is addressing the base of a floor-type milling machine. Komatsu provides an established route that can be analysed and improved, while WWR and Goimek test whether MaaS can help subcontractors prepare for unfamiliar work.

"Komatsu is an original equipment manufacturer, so it can either produce in-house or subcontract its production," explains project partner Alex Iglesias of Ideko. "We can see the performance of the platform where there is already a route for how to do it, and optimise that process through the platform. WWR and Goimek are subcontractors producing these large parts for the first time. They need the platform to help them maintain process stability, assure quality and keep production cost-efficient."

The need to prevent defects runs through all three demonstrators. REED's simulations, digital twins, monitoring systems and learning feedback loop are intended to identify risks before and during production. For a component that may take weeks to produce, quality cannot depend solely on detecting a problem during final inspection.

This connects REED's MaaS model with the principles of zero-defect manufacturing and its participation in the **4ZDM cluster**. "The cluster brings together different European projects working on zero-defect manufacturing using digital technologies, data and intelligent monitoring to prevent defects rather than correct them afterwards," says REED communication manager Vasiliki-Eleni Provopoulou of Netcompany. "This is particularly relevant for REED, where any errors or rework can significantly affect cost, materials, energy and production time." The collaboration also gives the consortium access to related work on process monitoring, data-driven quality assurance, AI and digital twins.

Capacity focus

Before REED can support a stable process, it must establish whether a manufacturer has both the equipment and the knowledge to deliver the component. Iglesias sees the shortage of this qualified capacity as a principal point at which bulky-part value chains break down. Large-part capability is concentrated among relatively few companies, and having suitable machinery does not necessarily mean that a supplier is prepared to accept the work.

"Even if they have the physical capacity, they can lack the knowledge," he says. "These big parts are connected to very specialised human knowledge about how to handle them, how to align them, how to process them and which cutting strategies to use. A company may be capable of doing it, but not dare to take on the work because, if it goes wrong, it is very costly."

Iglesias distinguishes between component-specific equipment and methods that can transfer between manufacturers. Digital twins for characterising the static and dynamic stiffness of a machine structure can be applied relatively widely. Workholding equipment is more specific because every component presents different dimensions and support requirements. "What can be transferred is the method of simulating, designing and aligning the parts," he explains. "Even if you are using different physical units, you are using the same procedure."

Process monitoring is harder to generalise because machine manufacturers and sensor integrators use different platforms, variables and units. Data from each system must be mapped into REED's shared data model so that measurements are defined and expressed consistently. Small batches also provide few repetitions from which the digital twins can learn, requiring the platform to apply knowledge from one component carefully to another that may be similar without being identical.

Iglesias regards differences in operator practice as another source of learning. If REED can capture alternative production routes, he thinks the platform can compare their cost, environmental performance and productivity. "We can learn which is the optimum route and use it the next time the part has to be produced," he says.

De la Peña argues that this learning model will only be acceptable if manufacturers retain control of sensitive information. REED is using a manufacturing data space in which information is exchanged under defined rules and only when required. "The data is very valuable," he says. "The client is going to ask for a piece, but the way of doing the piece is the valuable thing, and that is not going to be shared."

A supplier may expose its capabilities and submit a price, delivery time and quality offer without disclosing its computer numerical control (CNC) programme.

A responsive value network

REED's methodology follows an order from initial request to final delivery. A customer submits a request for quotation, and the platform checks the capacity and resources declared by onboarded manufacturers. "Two simulations run in parallel," consortium member Panagiotis Kourempanas explains.

"The process simulation hub feeds an environmental-footprint simulation, while the machine kinematics and machining-process twins validate feasibility. Both converge in a machine-learning component that generates production scenarios, followed by optimisation and selection of the best quotations on cost, time and environmental criteria."

If the customer accepts a quotation, the platform configures the participating network and initiates production. Three-dimensional process tracking, environmental monitoring and a factory-level KPI scoreboard show what is happening. If a threshold is breached, the system can alert the operator, amend the schedule or trigger a reconfiguration. Once the part is complete, production and lifecycle



information is assembled into a digital product passport and verified using blockchain before delivery.

Kourempanas describes digital twins at four connected levels. Machine-level twins simulate kinematic behaviour and inform feasibility and cutting decisions. Process-level twins compare sensor feedback with predictions to support interventions on vibration and tool wear. A visual process stream gives operators production-level oversight, while value-network models test the supply chain, logistics and production network under disruption scenarios.

"The validation is an ongoing cycle rather than a one-time check," he adds. "Operator feedback, measured deviations and the results of each production run feed a learning loop intended to improve the models used for future orders."

Hardware and software enter the platform as reusable services through standard interfaces. Cutting tools, vibration-control equipment and other devices retain their own functions but expose them in a form that REED can connect to the digital twins. Data-space connectors provide authentication and a common vocabulary, allowing most tools to be integrated without creating a unique connection each time.

Automated quotation and negotiation take place before an order is accepted. Scheduling, work preparation and monitoring follow afterwards. AI can analyse sensor data, KPI trends and operator feedback, but the customer and supplier retain responsibility for the industrial commitment.

Laboratory evidence

REED has developed its standalone tools and begun integrating them into the platform, while the industrial demonstration phase was beginning at the time of the interview. "We are replicating the use cases in the lab," Iglesias says, "trying to reproduce the manufacturing environment using the same machine, the same tools and the same strategies." The technologies had not yet been proven through complete production at the three sites.

"The results are very promising," he adds. "We have achieved good results in productivity increase, tool-life increase, dynamic-stiffness improvement and workpiece-alignment capacity."

One back-face milling tool more than tripled productivity in the replicated laboratory environment. Iglesias believes it could apply beyond the use case because back-face milling is common in large-part machining, but industrial validation remains the next test.

The team will first apply individual tools to selected operations at each site and use the results to improve their integration. By the end of the project, it aims to assist the complete production of a large component through the MaaS platform and compare it with the original process. The targets cover machining lead time, waste and carbon footprint, alongside machinery use, cost and productivity. Until the site work is complete, these remain targets rather than results.

Joining the network

"They must follow a structured onboarding and qualification process," de la Peña says. "They must register and verify their legal status, contracts and certifications, then provide their capabilities: supported materials, tolerances and quality levels. They must keep updating their actual availability, lead time and batch size." A quotation should therefore reflect the plant's current position rather than a static equipment list.

The provider must also connect to the manufacturing data space and enable the platform to exchange information with its own systems. Because production may start weeks after an order and take several more weeks, the supplier must report progress so the customer can follow the job.

Traceable process data, digital twins and the digital product passport are intended to make supplier qualification faster, but they do not replace engineering assurance. "What we can avoid are unnecessary trials and repeated searches for data," de la Peña says. "The information will be there, accessible when it is needed and trusted."

Blockchain verification is intended to show that the record has not been altered, giving engineers better evidence without replacing their judgement or removing customer accountability.

Consortium member Konstantinos Chisiridis (Exploitation Manager of the REED project) from Netcompany S.A. identifies machining subcontractors, tier-one and tier-two suppliers, and bulky-part OEMs as the principal users.

The consortium is defining the intellectual-property governance, operation, hosting and support needed for the integrated platform, with a subscription or annual licence as one possible model. Individual hardware and software components could also be sold separately, incorporated into their owners' portfolios or released as open source.

"The platform is the key result of the project," Chisiridis says, "but the idea is not only to exploit the platform. We may create service bundles from selected software components for very specific needs of end users. This would allow companies to adopt selected REED services without taking the full environment at once."

The human remains in the loop

REED began with an ambition to digitise manufacturing from end to end. In practice, variation among machines, sensors and operator methods makes complete automation unrealistic within the project. Data from each monitoring system must be mapped into the platform's shared data model so that measurements are expressed consistently. CNC programming is more difficult to automate. REED originally considered generating manufacturing routes automatically, but a large-

part programme contains too many geometry, machine, setup and process decisions for that to be dependable today. The revised approach asks the MaaS user to provide an initial CNC programme, which the platform uses as the basis for calculating alternatives.

Iglesias says the project has adjusted its approach accordingly. "We did not want to get stuck on these side problems," he says. "We thought it was better to leave some things unsolved or not digitised for the moment. In summary, the platform's decision process still needs human involvement."

REED is therefore formalising the evidence around operator experience so that people can make better-supported decisions and useful knowledge is retained for the next order.

The route to a live service

Evidence will emerge incrementally from the three industrial sites. Partial production trials will test individual tools before later demonstrations connect them across a complete component. Iglesias expects progress in process monitoring, sustainability measurement, cost efficiency and productivity, but says the final results must show how closely REED has met its technical commitments.

Commercial work will develop alongside the trials. "The idea is to use these three demonstrators as potential early adopters when the technical and business benefits are demonstrated by the end of the project," Chisiridis says. "That would provide a basis for paid agreements and for approaching manufacturers outside the REED consortium with similar problems. Initial offers may consist of individual tools or service bundles rather than the complete platform, since the MaaS business model of REED is quite nimble and flexible."

De la Peña expects tools and improved cutting technologies to reach the market and sees scope to reduce the time consumed by alignment and clamping. "There are not so many producers in the market; they are few, but they have a lot of potential," he adds. "If we offer them the possibility to provide these services to new customers who need their capabilities, that would be a perfect result."

The environmental leverage could be considerable because a bulky component can embody tonnes of carbon, substantial material and days of machine time. Better planning, monitoring and traceability should reduce both commercial exposure and avoidable environmental cost.

Europe, de la Peña argues, already has companies that lead this specialised market. REED will have made a lasting difference if it allows them to accept unfamiliar work with better evidence, share only the information required for the transaction and learn systematically from production without surrendering the expertise that gives them their advantage.

Customers need confidence in delivery, while suppliers need confidence that participation will not expose the methods on which their businesses depend. "Data privacy is one of the key pain points for customers," Chisiridis says. "They have control of their data and no sensitive information is disclosed. That protection must be matched by evidence that a newly configured value network can deliver the required part."

If REED can establish both, MaaS for bulky parts becomes less a promise of frictionless automation than a practical way to put scarce European capability to wider use.

PROJECT INFORMATION

Project Title
REED – REsponsive and REconfigurable value network for the manufacture of bulky parts

Project Objective
The main objective of REED project will be the development of a Manufacturing as a Service (MaaS) platform that will provide the enabling technologies, equipment and services for the manufacturing of bulky parts of the capital goods sector, assuring the quality of the produced components while maximising productivity and minimising environmental impact.

Project Duration and Timing
39 months. Starting date: Oct. 2024
End date: Dec. 2027

Project Funding
Granting authority: European Health and Digital Executive Agency. Grant amount (EU Contribution): € 5 295 168,13

Project Partners
I3B- Instituto Ibermatica De Innovacion Sl (Project & Technical Coordinator); INTRA- Netcompany SA; SIN- Sintef AS; IDK- Ideko Scoop; LMS- Panepistimio Patron; TUDA- Technische Universitat Darmstadt; SMT- Smarttech Industry Software Solutions; 2LC- 2.-0 Lca Consultants Aps; MOD- Module Works Gmbh; SOR- Soraluce Coop; MAT- Dr Matzat & Co Gmbh Spann Und Fertigungstechnik; KOM- Komatsu Uk Limited; WWR- Maschinenfabrik Wustwillenroth Gmbh; GOI- Goimek

MAIN CONTACT



Name
Vasiliki-Eleni Provopoulou
Netcompany
eleni.provopoulou@netcompany.com

Pedro Maria de la Peña Tejada
Instituto Ibermática de Innovación SL – Ayesa
pmdelapena@ayesa.com



Email
info@reed-project.eu

Web address
https://reed-eu.eu/

REED

The REED project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement **No. 101178405.**

INSIGHT OPINION

4ZDM
Zero Defects Manufacturing

Zero-defect manufacturing: Europe must build on its investment

Europe should launch a new funding call for Zero-Defect Manufacturing (ZDM). Successive investments have built considerable expertise, developed promising technologies and produced industrial demonstrations that work well on the factory floor. That investment should not be wasted, and continued efforts should now be made to turn real progress into sustained manufacturing capability.

Through 4ZDM, a group of researchers, projects and leading manufacturing companies, Europe has a community able to bring the evidence together, connect industrial needs with research expertise and help shape an ambitious funding agenda focused on delivering practical, commercial outcomes.

The case for renewed investment starts with a persistent challenge: quality problems are often discovered after costs have accumulated. A defective component may have consumed materials, energy, machine time and labour before final inspection identifies the fault. ZDM uses production data to anticipate problems, guide intervention and get products 'right first time', avoiding costs while improving competitiveness and resource efficiency.

The technology underpinning this opportunity is increasingly strong. In-line sensing and non-destructive inspection reveal what is happening during production. Machine learning and virtual metrology help estimate quality from process data. Predictive models and digital twins support decisions about emerging deviations, while adaptive control and repair strategies create opportunities to recover value before a component becomes scrap. The challenge is to combine these capabilities into dependable systems that work within industrial constraints.

European research has helped build those foundations and advance the state of the art. ZDMP pursued a shared platform connecting manufacturing assets and quality applications. Work involving ZDMP and ZDZW also helped initiate common ZDM principles and requirements through CEN-CENELEC. Such collaboration has helped address manufacturers' need for technologies that communicate, integrate with existing equipment and support consistent approaches to quality. Each project has built on what came before and delivered results that others can use.

The FLASH-COMP and TURBO projects illustrate the direction of travel for industrial deployment. Though not ready for full integration at all levels of complexity on the factory floor, their industrial validations at Azimut (see page 83) and at Siemens Gamesa Renewable Energy brought together physical sensing, visual monitoring, data analysis and predictive simulation to support decisions during composite manufacturing. This is a practical expression of ZDM, one in which understanding a process while there is still time to influence its outcome delivers valuable real-time data that operators can use to make a difference. It also exposes the next challenge, however, which is turning a successful demonstration into repeatable implementation across variable production environments.

A new call should make that transition a central topic, with industrial partners helping define the research agenda, identifying costly quality problems, setting performance requirements and committing the production access, data and people needed to address them. Use cases should demonstrate value at industrial scale, under realistic variation in materials, operating conditions and product mix. Convenient trials of easily measured defects cannot establish whether a solution will survive the demands of daily production.

Success should be measured against agreed baselines: first-pass yield, scrap, rework, throughput, resource consumption and total deployment cost. Demonstrations must also establish how operators understand recommendations, how systems respond when predictions are uncertain and who takes responsibility for intervention. Human expertise, maintainability and trust are essential to deployment.

Crucially, the route beyond the grant must be credible and funded. Proposals should identify who will own, operate, maintain and update the resulting systems; how integration and workforce training will be paid for; and what investment will support continued operation after project funding ends. A strong exploitation plan, supported by good communications, needs named responsibilities, committed resources and a route to wider adoption. Follow-on finance and replication support should help promising results cross the gap between demonstration and routine use.

The 4ZDM Cluster is well placed to give this agenda a collective voice, by bringing together projects across successive European research programmes, and being connected to ManuFuture. It offers a home for sharing evidence, comparing approaches and forging collaborations. It can bring together manufacturers, technology providers and researchers to identify unresolved barriers and translate them into priorities for future calls. And through its industry involvement, it offers an honest account of what has worked, what remains difficult and what deployment requires.

Europe has already invested heavily in ZDM, and a new funding round can now build on that investment. Any resulting funding should demand stronger industrial ownership and support deployment that lasts. 4ZDM can help shape that call and mobilise the expertise to deliver it. The ambition now must be to make the benefits of ZDM endure on Europe's factory floor.

SMARTER, MORE FLEXIBLE FACTORIES

ONE4ALL: Reconfiguring manufacturing around people

Manufacturers need to respond to disruption, changing demand and shorter production runs, yet many – particularly SMEs – rely on rigid equipment, manual planning and limited data that make rapid change difficult. ONE4ALL is addressing this challenge by combining reconfigurable collaborative robotics with digital twins, intelligent data management and decision support. Together, this creates a framework that helps factories anticipate change, adapt production and keep operators in control. Project coordinator **Doru Stefan Andreiana** explains how the technology works and how two contrasting demonstrators are bringing this novel approach to life.

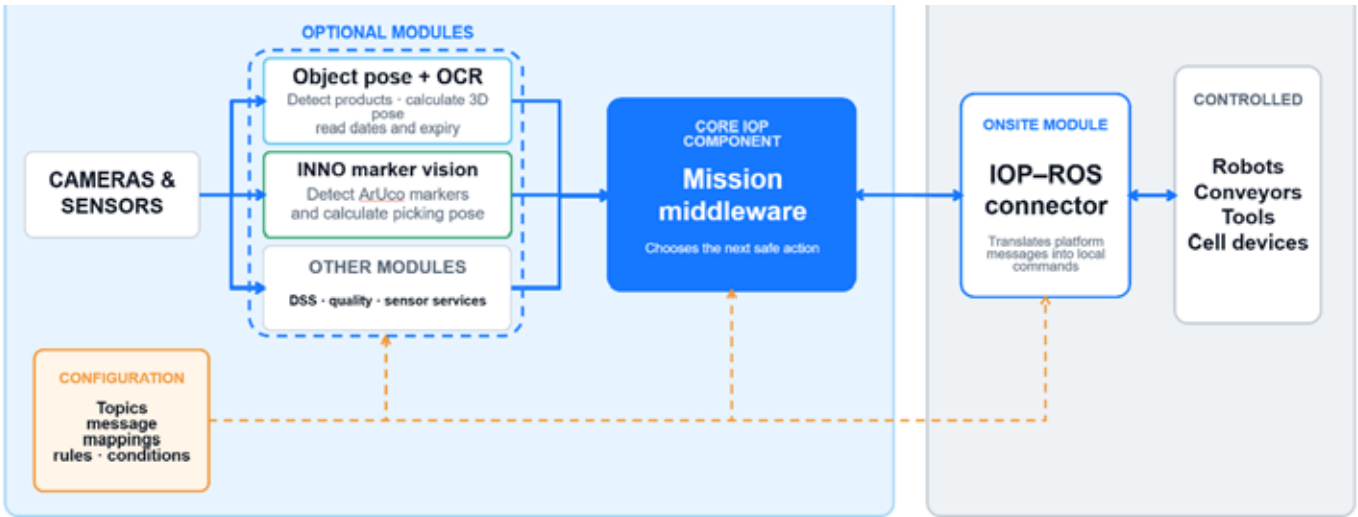
The shock of COVID-19 exposed a common weakness in manufacturing that exists even outside such extreme challenges faced in a pandemic. Production systems designed for stability and scale can be highly efficient when demand behaves as expected, but they are far less capable of changing direction when demand shifts unexpectedly. In the pandemic, new products were suddenly needed, established orders were disrupted, and factories had to reconfigure at a speed that conventional production planning had rarely anticipated.

That experience provided the inspiration for ONE4ALL, a Horizon Europe project seeking to make manufacturing more resilient by enabling plants to adapt quickly when needed, both physically and digitally. Its particular focus is on SMEs where the effects of disruption can be more severe and the resources available for transformation more limited.

"One of the aims of the project is to enable the manufacturing industry, especially SMEs, to reconfigure and adapt to changes in demand, production or anything that cannot be predicted," explains Doru Stefan Andreiana, ONE4ALL project coordinator and a robotic and mechatronic engineer at IDENER, the Spanish research SME coordinating the project. Specialising in computational solutions for manufacturing, AI applications and modelling, Andreiana and his team have helped steer the project from its original proposal through the integration of its technologies and their demonstration in industry.

Andreiana points out that the response ONE4ALL proposes is deliberately broader than automation. "A factory cannot become genuinely adaptable simply by installing a robot, nor can a digital model create flexibility if the machinery on the floor cannot change," he explains.

Reconfigurable cyberphysical production module (RCPM) software framework



The project therefore connects three layers: modular physical equipment that can be redeployed, digital systems that understand and anticipate production, and a human-centred decision-support framework that keeps operators in control while measuring the broader sustainability of change.

"First, you need reconfigurability at the mechanical level, and that needs to be supported by the digital part," Andreiana continues. "Then, as a foundation for everything, you have the Industry 5.0 approach: ensuring that the operators are always in the loop. Instead of replacing them with these technologies, we support them to do their work more efficiently."

Flexibility on the factory floor

The physical expression of this idea is what the project calls a reconfigurable cyber-physical production module, or RCPM. In practice, it is a mobile reconfigurable robot, built from a wheeled base, a robotic arm and interchangeable tools. Its mobility allows it to reinforce different processes; its modular grippers allow it to handle different products; and its cameras, sensors and control software connect the physical machine to a company's wider digital architecture.

That is an important distinction, as Andreiana explains: "Conventional automation is often engineered around a stable environment in which every object arrives in a known position and every action is repeated. ONE4ALL is aimed at more dynamic factories, where product type, orientation and production priorities may change. Its vision system acts as the robot's eyes, identifying an object and its position so the control system can select the appropriate tool and movement.

"It is not just automation, but something dynamic," Andreiana continues. "The mobile robot can move to different parts of the factory, different grippers can be used to handle different objects, and the vision system enables it to work in a human and dynamic environment."

For an SME, this potentially lowers the threshold for flexible automation. The host factory does not need to begin with a fully digitised production line, although the workspace must be assessed carefully. There must be sufficient room for the robot to operate, while its speed and separation from people must meet relevant safety requirements on site. A new application may also require a purpose-designed gripper. The principle, however, is to adapt a reusable module to the plant rather than rebuild the plant around a single fixed machine.

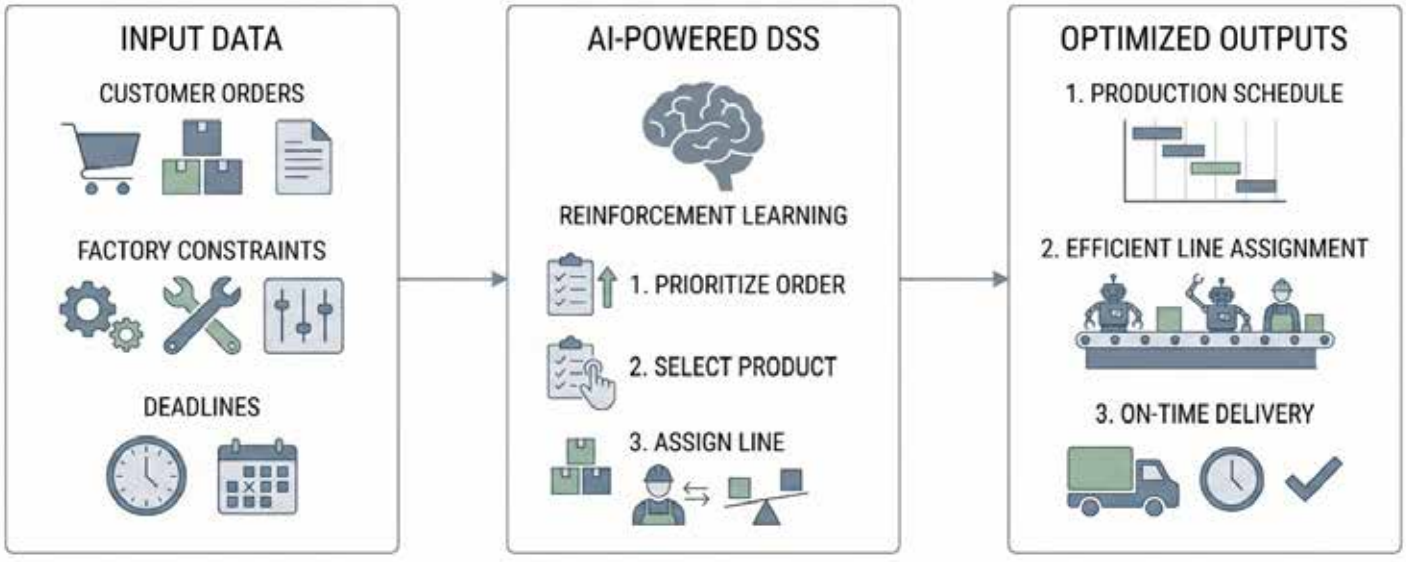
From factory data to usable decisions

Physical flexibility only becomes valuable when a manufacturer knows what to change, when to change it and what the consequences are likely to be – and that needs data, and ONE4ALL's digital layer is designed to provide that intelligence. At its centre is an intelligent orchestration platform, or IOP, which connects the robot, sensors and project digital services with any manufacturing execution, enterprise resource planning or other IT systems already used by the company.

The platform's data-fusion pipeline can ingest different kinds of information arriving through different communication channels. Andreiana describes the IOP as the messenger: "It ensures that information can flow between otherwise separate machines and services," he says. "Digital twins then use that information to represent production and explore "what-if" scenarios, while a distributed decision-support system turns the resulting insight into recommendations."

Consider the planning problem faced by many smaller manufacturers. Orders and jobs may still be scheduled manually, sometimes on paper. Creating a weekly plan takes time; inserting an urgent, short-run order can mean rebuilding it. ONE4ALL's decision-support system can generate and update a schedule while considering deadlines, available equipment, materials, energy use and business priorities. The

Conceptual diagram of Decision Support System for dynamic scheduling of orders and resources



planner checks, adjusts and approves the recommendation rather than surrendering control to it.

"If a new order comes up that you did not plan, you have to change your schedule, which takes a lot of time," Andreiana says. "The decision-support system can adapt the schedule while maintaining deadlines and KPIs. But the decision is still up to the workers and managers. We assess and show the expected impact; they decide whether to follow it."

The digital twin extends this responsiveness beyond events already encountered. By simulating unexpected small batches, urgent demand or the loss of a production resource, it can create scenarios on which the decision-support system is trained. "If a similar disruption occurs in the factory, the system is better prepared to propose a workable response," adds Andreiana.

The data gap

This approach has exposed one of the most important barriers to industrial digitalisation: the companies that stand to gain most from better decisions may have the least data with which to build them. ONE4ALL's demonstrators were not highly digitised, a situation Andreiana says is common among SMEs, especially hand-labour-intensive types. Records may be incomplete, inconsistent or held on paper; sensors may not exist; and real-time data may be unavailable.

This does not make digital tools impossible, but it changes what they can deliver. Where sufficient quality data exists, a digital twin can learn a model of the factory and update as conditions change. Where it does not, the project draws on another valuable source: the people who already understand the process.

"Even if the company does not have data, it most probably has operators," Andreiana observes. "They know how the machines work, how much time it takes to make a product and how the factory works and through questionnaires and structured knowledge capture, that expertise can be used to construct an initial model.

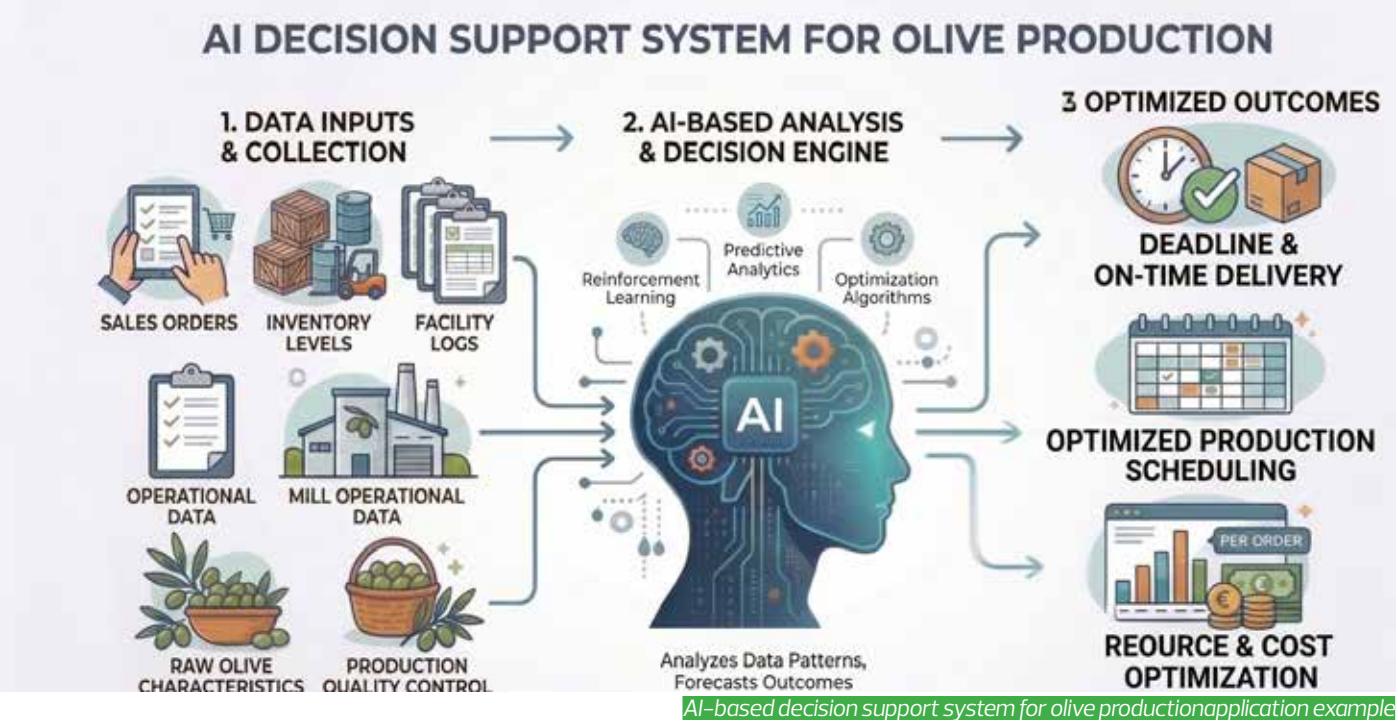
"Paper-based orders can also be digitised to support scheduling," he continues. "The result lacks the full real-time adaptability of a data-driven twin, but it offers an attainable starting point rather than making perfect digital maturity a precondition for progress."

This lesson is particularly significant for SMEs. The barrier is not only the cost of technology, but the quality of the information infrastructure, the skills available to use it and the organisational readiness to change. ONE4ALL therefore incorporates digital-maturity and sustainability assessments at the beginning, middle and end of the project. These establish where a company is, identify obstacles and create a route towards the level of digitalisation it actually needs.

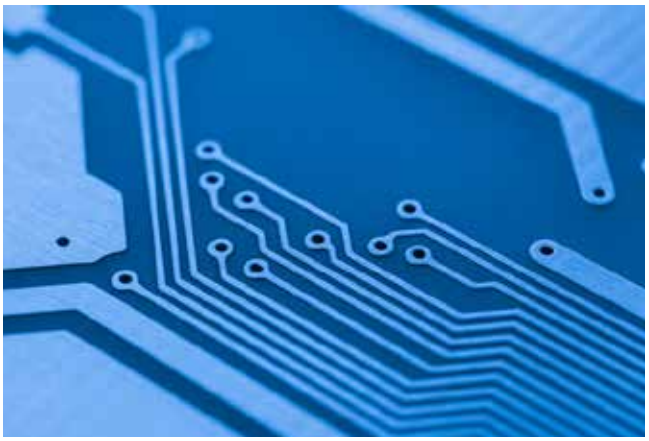
Training is linked to that assessment and to the design of the technology itself. Workforce needs and difficulties are fed back to developers, while guidance helps operators use the robot and digital tools. The assessment remains valuable even if a business ultimately chooses technologies outside the ONE4ALL portfolio. As Andreiana puts it: "It is always useful to know where you are, what you should do to reach the next level and where you would like to be in a few years."

Trust, safety and interoperability

Keeping people at the centre also means addressing the conditions under which they can trust and safely use the



AI-based decision support system for olive productionapplication example



system. ONE4ALL identified worker-safety requirements at the start of the project and incorporated them into the robot's mechanical design, software and control. Before commercial deployment alongside operators, the physical system will need further optimisation, higher technology readiness and satisfaction of all application-specific regulatory and standardisation requirements.

The project has treated cybersecurity, interoperability and trustworthy AI as cross-cutting concerns. Access controls and security mechanisms govern data use within the IOP, while its flexible data pipeline allows different services and message types to connect. The robot's own components – including its arm, controller and vision system – must also interoperate reliably; the team used the open-source Robot Operating System 2, ROS2, to help integrate them. The cameras and sensors themselves are commercially available, while the project's innovation lies chiefly in the AI-based software models that interpret images and locate objects.

These issues are not peripheral. A modular production system only becomes reusable if its components can communicate, and human oversight only becomes meaningful if users can understand the recommendation and retain authority over the action. ONE4ALL's emphasis on open approaches, user-facing information and decision support reflects the practical reality that adoption depends as much on confidence and integration as technical performance.

Practical pathways

As ONE4ALL approaches completion, the technologies are moving towards different forms of exploitation rather than a single, finished "one size fits all" package. The digital twin, decision-support system and maturity assessments are the most immediately transferable. They can operate independently of the robotic module and, subject to the availability of data, be adapted across sectors. The physical module is more likely to reach market first through tailored applications, while further work raises speed, readiness and compliance for broader deployment.

That separation does not diminish the value of the integrated project. On the contrary, the demonstrations show that manufacturers need an entry point suited to their circumstances. One business may begin by digitising orders and improving scheduling; another may model its energy and material use; a third may deploy a RCPM to a slow, hazardous or ergonomically difficult process. The assessment framework can show what is feasible now and what capabilities need to be built next.

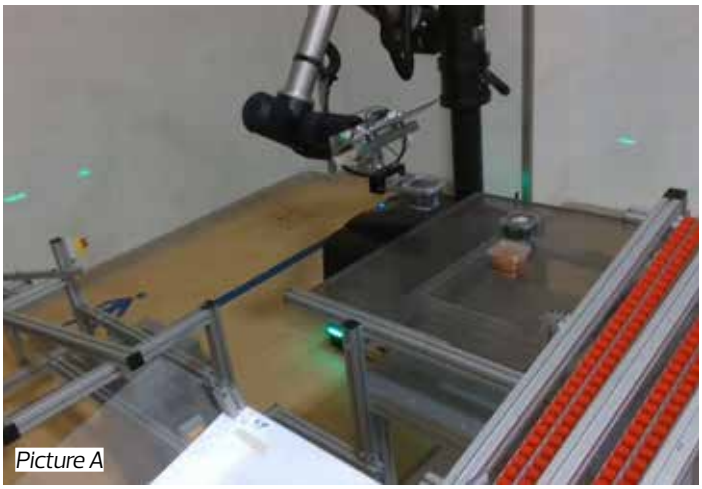
Success will ultimately be measured in responsiveness: how quickly a factory can understand a change, evaluate its options and act without undermining safety, sustainability or human control. "At the digital level, we increase response speed because scheduling can be done faster and adapted in real time," Andreiana says. "With the digital twin, you can also predict the scenario. At the physical level, the robot enables you to reinforce a process when it is needed."

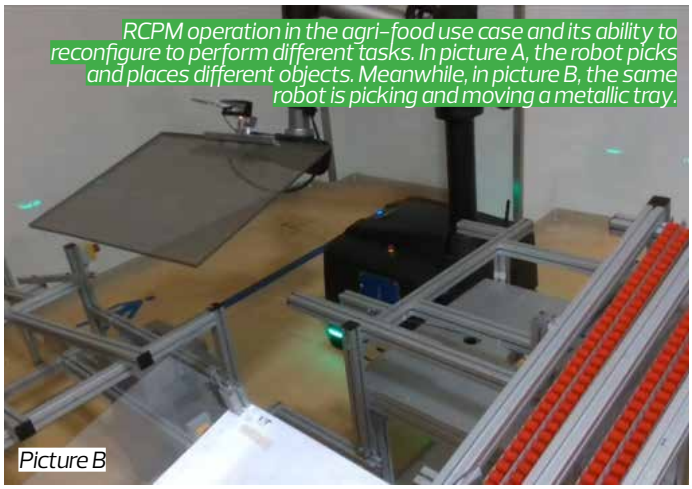
There is a final realism running through the ONE4ALL story. Moving technology from a controlled test bed into a working factory is difficult. Machines must meet real production rates; separate components must operate as one; data is rarely as clean as anticipated; and installation cannot be allowed to jeopardise customer deadlines. Andreiana describes this scale-up and integration as the project's principal challenge. Yet it is also where the most transferable insight has been generated.

For SMEs in particular, resilience will not arrive through a dramatic leap from manual production to an autonomous factory. It is more likely to be built in stages: capturing knowledge, improving data, supporting a planner, adding visibility, and then placing flexible automation where it creates clear value. ONE4ALL's novel contribution is to connect those stages while retaining a straightforward principle – that technology should make the manufacturer more adaptable by supporting its people.

Reconfigurable cyberphysical production module application example




Picture A


Picture B

RCPM operation in the agri-food use case and its ability to reconfigure to perform different tasks. In picture A, the robot picks and places different objects. Meanwhile, in picture B, the same robot is picking and moving a metallic tray.

Olives, variation and urgent orders

ONE4ALL's first industrial demonstration has been taking place at Madama Oliva, an Italian family business producing table olives and related products. It is a revealing setting because agri-food manufacturing combines labour-intensive operations, natural product variation, multiple formats and a large SME base. It also demonstrates why adaptation requires more than a single technology.

At the beginning of the process, olives must be separated according to quality. The project initially explored using the collaborative robot, but water in the environment and the speed required made that approach unsuitable. "Instead, we implemented a dedicated sorting machine reinforced through AI-based plugins. We developed an AI-based vision system to identify and reject defective olives, and a predictive Quality Statistical Supervision System (QSSS) as a decision support system.," Andreiana explains. "We improved the vision software so it could recognise different olive types and defects, added monitoring to give operators quality statistics, and enabled them to adjust the machine configuration on-the-fly, supported through predictive tools."

Despite this automation, workers are still very much involved on the production line with operators still checking the results, dealing with imperfections and configuring the process. More importantly, the decision to change technical direction is itself a useful demonstration of human-centred engineering. "The technology was adapted to the reality of the task rather than forcing the original concept into an unsuitable environment," Andreiana points out.

The project's more ambitious robotic demonstration sits at the end of the line. Products arrive in different jars, pouches and boxes, in varying positions. Using its vision system, the RCPM identifies each item, checks information such as the label and expiry date, rejects anything incorrect, and moves acceptable products onwards for palletisation. Once a tray is empty, the robot changes its handling configuration to remove it, then prepares for the next set of products.

This physical flexibility is connected to the production picture maintained by the IOP and the schedules proposed by the decision-support system. "Madama Oliva's predictable domestic production is often made for stock, but export orders can arrive unexpectedly, in uncertain batch sizes and with short deadlines," says Andreiana. "This is precisely the kind of high-priority variation ONE4ALL was designed to address."

The planners' early response to the digital scheduling system has been positive, Andreiana reports, because order scheduling previously completed manually can now be generated, checked and adapted in real time, while considering multiple constraints and objectives. The robotic system remains a demonstrator on a replica line, however. Its modularity, perception and ability to work in a changing environment have been proven, but the overall performance and speeds still has room of improvement.

"We have the technology, and we know that it works; now we need to scale it up," he says. "The value is that it is modular, configurable, can work around people and can work in a dynamic environment. We now need to optimise it, which is work to be done after the project ends."

A value case in pharmaceuticals

The second demonstrator, hosted in a pharmaceutical training laboratory at InnoGlobal in Ireland, explores a different balance of time, safety and productivity. Because the facility replicates pharmaceutical production rather than processing commercial orders, the team uses simulation to extrapolate how the tools could operate at factory scale.

In controlled pharmaceutical environments, moving between products can demand intensive cleaning to prevent cross-contamination. Workers may also be exposed to hazardous processes, and entering and leaving a clean workspace can itself consume significant time. Andreiana gives the example of a five-minute task that might require around half an hour of preparation to enter and another half an hour to leave if the

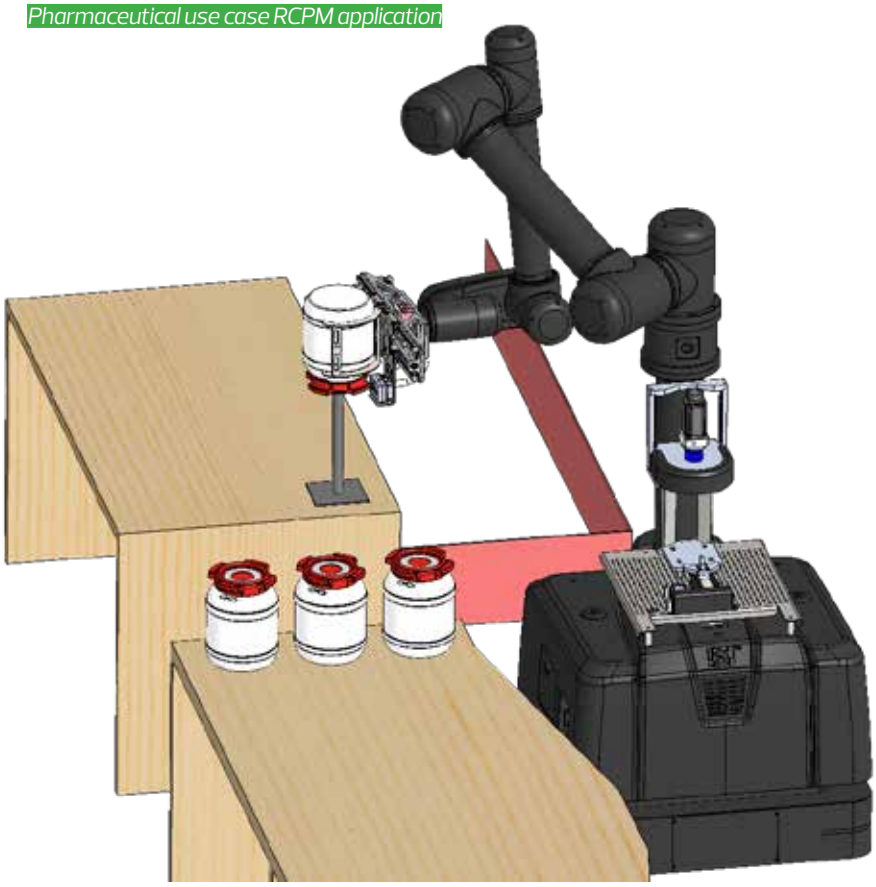
production area is one that processes high-potency active ingredients.

Here, the RCPM's value is not primarily raw speed but its autonomy in a dynamic environment and its reconfigurability. It can remain inside the controlled area, handle and mix a container of pharmaceutical powder, move it to the next process step and support quality checking. The digital twin monitors factors such as energy and resource use, while the decision-support system explores where a mobile RCPM would add the most value across a scaled production plant.

"Where should you assign the RCPM so workers do not have to lose an hour going in and out, and so they do not have to be exposed to an activity that can be dangerous?" Andreiana asks. The answer may change according to the production campaign, urgency and cleaning requirements. In this setting, the project's reconfigurability becomes a way of deploying automation where it best protects people and removes avoidable delay.

The contrast between the two pilots is instructive. Beyond reconfigurability, the agri-food use case spotlights a hand-labour-intensive and variable environment demanding throughput, high adaptability and human-robot collaboration; meanwhile, the pharmaceutical setting places a higher value on controlled handling, reduced exposure and avoiding repeated cleanroom entry. A technology's readiness, in other words, cannot be judged by capability alone. Its value depends on the process into which it is introduced.

Pharmaceutical use case RCPM application



PROJECT INFORMATION

Project Title
ONE4ALL – Agile and modular cyber-physical technologies supported by data-driven digital tools to reinforce manufacturing resilience

Project Objective
ONE4ALL aims to reinforce the resilience and flexibility of the manufacturing lines, SMEs-focused, through innovative modular and reconfigurable technologies centred on humans and sustainability.

Project Duration and Timing
48 months

Project Funding
HORIZON-CL4-2021-TWIN-TRANSITION-01-03 (GA ID 101091877)

Project Partners
IDENER RESEARCH & DEVELOPMENT AIE
INNOPHARMA RESEARCH LIMITED
CRIT CENTRO DI RICERCA E INNOVAZIONE TECNOLOGICA SRL
EXELISIS IKE
SYDDANSK UNIVERSITET
INNOBOTICS SRL.
MADAMA OLIVA SPA
HOLOSS – HOLISTIC AND ONTOLOGICAL SOLUTIONS FOR SUSTAINABILITY, LDA.
TECHNISCHE UNIVERSITAT DORTMUND
KARLSRUHER INSTITUT FUER TECHNOLOGIE

MAIN CONTACT

Name
Doru Stefan Andreiana



Email
doru.stefan@idener.ai

Web address
www.idener.ai

This project has received funding from the European Union's Horizon Europe research and innovation program under the grant agreement **No 101091877**.

SMARTER, MORE FLEXIBLE FACTORIES

AIRISE: Turning AI into a manufacturing mission

For many smaller manufacturers, the question is not whether artificial intelligence sounds promising, but how to turn it into a safe, affordable and useful application on the factory floor. EU-funded AIRISE has created a guided route for companies to follow from readiness assessment to working demonstration. This is supported by an ecosystem that brings together manufacturing knowledge, AI expertise and commercial solutions. Project coordinator **Harry Bikas** and technical manager **Ulrich Thombansen** explain how this approach will continue to deliver the benefits of AI to SMEs long into the future.



Artificial intelligence has become almost impossible for manufacturers to ignore. Yet the leap from asking Copilot to help organise your emails to implementing a dependable application controlling quality, scheduling work or monitoring a machine on the factory floor is enormous, particularly for SMEs without the data scientists and software developers to help, or the time to experiment.

Ulrich Thombansen, AIRISE technical manager at Fraunhofer Institute for Laser Technology ILT, describes the typical starting point for most SMEs simply: "People hear about AI and think, 'That's cool,' but they come in with one question: 'What's in it for me?' They want to use it because they want to get better, but they don't have a clue about the how."

AIRISE – Artificial Intelligence in Manufacturing for Sustainable Applications at SMEs – was created to answer that 'how' question in practical terms. Coordinated by the Laboratory for Manufacturing Systems and Automation (LMS) at the University of Patras, the Horizon Europe initiative does not begin by inventing an AI model and searching for somewhere to use it. It begins with that initial pain point most SME

manufacturers face – what can AI do for us and how can we implement it?

"It is a strange project in the sense that it does not really develop new technology on its own merit," says project coordinator Harry Bikas of LMS. "It is here to support SMEs in manufacturing to get started with AI solutions on their shop floor."

"In that sense, we don't like to call AIRISE a project," adds Thombansen. "We like to think of it as a mission to show SMEs how to use AI and provide the support they need, so it can deliver benefits to their manufacturing processes, like process efficiency, decision support, and waste, energy and cost savings, for example."

It has addressed this aim in practical terms through a series of open calls, in which companies have described their challenges, how they believed AI might help and why they needed outside expertise. AIRISE then combined financial assistance with tailored technical support. "Companies provide data and know-how," Bikas explains. "We then train algorithms with their data and create a customised AI solution that solves their problem and this is then deployed in their environment, on their shop floor."

Beyond the AI headlines

There is a significant distinction between general-purpose AI and industrial AI – and this is crucial for AIRISE. A business can easily find a language model to support marketing or customer service, but a manufacturer trying to detect vibration during milling, optimise a food-production process or translate design information into a bill of materials needs a specialised combination of process knowledge, usable data, sensors, computing infrastructure and algorithms.

"You can't go to Amazon and ask, 'What do I need to do to make my packaging machine more intelligent?"' Thombansen says. "You need to know about that manufacturing challenge, and then you need to know about AI and how it can help meet that challenge. Otherwise, you cannot do this analysis."

That combination is the core competence that AIRISE has assembled. Its experts understand manufacturing at process and production level, but crucially they can also draw on expertise in AI, data, sensing, cybersecurity, business and human factors.

AIRISE turns that breadth of expertise into a single, coordinated offering for the SME. Rather than expecting a manufacturer to identify the right algorithm, technology supplier or data specialist, it first translates the company's operational problem into a clearly defined AI challenge. It can then assemble the combination of specialists and services appropriate to that particular case.

"We have specialised solutions that tackle specific problems for each SME that comes in," Bikas explains. "That is where we bring value: understanding their problem, breaking it down, seeing what AI can do for them, and then creating a solution around that."

As an entry point into industrial AI rather than simply another technology provider, AIRISE helps the manufacturer understand what is feasible, brings together the capabilities required to demonstrate it and creates a credible route towards implementation.

The tailored nature of that route also determines how the important question of cost must be approached. An intervention might require little more than a low-cost computer running a Python programme, or it might demand a complete edge installation. Before an SME can calculate return on that investment, it must define the problem, discover what data exists and understand the shape of a credible solution.

"The ROI discussion is so far away from the starting point that companies come in with," adds Thombansen. "They come in having a vague idea of where AI could help and we go from there, rather than starting with cost."

The value of data

Data is perhaps the biggest requirement – AI cannot work without it. Many companies AIRISE has worked with have plenty of data, but often they do not know whether it is complete, consistent or even relevant. Information may be recorded in spreadsheets, held on paper or locked inside the experienced heads of operators. Some factories remain dependent on manually operated equipment and must begin by installing sensors or digitising records.

Production and engineering data can also be commercially sensitive, and many manufacturers are understandably reluctant to expose valuable process knowledge to an external, publicly available AI tool and may prefer locally controlled models. AIRISE therefore treats data availability, quality, ownership, infrastructure and security as fundamental parts of the manufacturing problem.

AI navigation pathway

AIRISE's response to all these challenges is an AI navigation pathway that turns an AI ambition into manageable steps. Even before applying for deeper support on application, cost analysis and data acquisition, an SME can use the project's readiness assessment tool, either independently or complete it as a guided interview.

"Our assessment looks at how the company is structured and digitised, the competence available internally, its data and infrastructure, the likely business impact and the resources required," explains Thombansen. "It also begins to explore how a proposed application could affect people and their work.

"This gives us a snapshot of where the company is in terms



of how educated they are and how well equipped they are," Thombansen continues. "That starting point matters because apparently simple ideas may need to be narrowed, while some businesses must strengthen their digital foundations before an AI experiment becomes viable."

From that point, the assessment determines which capabilities the SME needs and in what sequence.

The work then moves through three distinct phases: initiation, implementation and validation. At the outset, the project team working with the company establishes the problem, approach and intended KPIs. During implementation, they develop the application together, while validation takes the solution into the SME's premises so that its technical and business value can be evaluated under realistic conditions.

The degree of hands-on support changes with each manufacturer. Where a company has capable people, AIRISE trains them and guides the work. Where it lacks those resources, the project can undertake more of the development while explaining what is being done so the SME can own the next step.

"We don't want to do turnkey solutions and just throw them over the fence and run away," says Thombansen. "It is really about skilling companies up, giving them the right tools, recommending the right steps and enabling them, as much as possible, to do the job themselves."

Proof through application

AIRISE has supported 75 SMEs across a deliberately broad range of applications so far. These include AI-assisted design, vision-based defect detection, process monitoring and control, predictive maintenance, production planning, shop-floor scheduling, logistics and resource optimisation.

"The common thread in all this work is not a particular algorithm, but a specific manufacturing problem with an outcome the company values," says Bikas.

One machine shop, for example, used a couple of inexpensive cameras to identify machine states, part cycles, idle periods and the causes of lost time. The resulting information fed into AI-assisted adaptive planning across its milling and turning equipment. This gave operators responsible for several machines much better visibility of where attention was needed.

Bikas explains: "The company almost doubled the utilisation of its machines and estimated that, even if it had financed the work itself, the investment could have been recovered in around a month. We were not expecting that, and they were not expecting that."

Other interventions have used cameras to make in-line quality decisions; adjusted process parameters in food production to reduce water consumption; detected unwanted vibrations during milling so the process could be adapted; and helped designers and planners explore alternatives. In job shops, scheduling tools can connect incoming orders with available machines and personnel, adapting when materials, staffing or equipment availability changes.

"Even a simple AI indicator can improve working life by directing an operator towards a part that may require attention instead of forcing that person to inspect everything continuously," adds Thombansen.

Bikas also recalls a familiar quality-control challenge becoming difficult because the company lacked consistent data. Conversely, a seemingly complex application that interpreted designs and automatically created bills of materials progressed smoothly because the SME worked in a structured way and could provide good information.

"AI is not a magic box where you put problems in and get solutions out," he reflects. "Let's see where you are before you start implementing."

Confidence through experience

The first implementation is also critical for building confidence among the people who will use the technology.

"They need to see and to understand what AI can do for them in practical terms," Bikas says. "Get the first implementation done and let them experience it first-hand, not something they read about, but something they see directly on their shop floor. Then the company can decide whether they want to invest further."

That experience can also bring questions to the surface. Thombansen says companies may not always state explicitly that employees fear losing their jobs, but "we often feel that worry, so a working demonstration makes a more honest discussion possible: what the technology will do, whether people will retain authority, how roles may change, what benefits are expected and what side effects the organisation must manage.

"Once decided, every stakeholder needs to understand what is going to happen," Thombansen continues. "You need an open discussion about the expected benefits and the effects on the company.

"Adoption also depends on keeping talent in the company," he concludes. "An SME that trains an enthusiastic employee must give that person time and support as well as a reason to stay with the company once their expertise becomes more valuable in the employment market. Technology adoption is therefore also a process of skills development alongside organisational change."

Ethics, explainability and conformity

Demonstrating that an AI application can work and create business value is only part of the AIRISE journey. As it moves towards implementation on the factory floor, the focus must widen from what the technology can do to how it will make decisions, affect people and operate within the manufacturer's existing structure. For Bikas and Thombansen, this brings three closely connected considerations into play – ethics, explainability and conformity – which need to be at the heart of planning from the outset. Together, they ensure that a solution will not only perform its intended task, but do so safely, transparently and responsibly, delivering value for the business while respecting the people whose work it will support or change.

AIRISE integrates ethics from the beginning rather than treating it as a final approval gate. Structured checks ask whether an application affects jobs or workplaces, which data it requires, whether any information is personal and whether it can be anonymised without diminishing the result. These questions are revisited after implementation and again during validation, allowing risks and recommended actions to evolve with the solution.



Explainability is especially important when an algorithm influences a physical process. "The solution or suggestion needs to be understandable by the operator, by the user," Thombansen explains. "The people designing it need to think about how to explain why it comes up with this suggestion".

"If a suggestion is not explainable, it may be tolerable if a person reviews it, but there is then a risk that operators will begin to ignore it," he continues. "AI in manufacturing has completely different rules from the open AI world. We cannot afford failures. We cannot afford a model freaking out and doing whatever it wants."

Conformity is equally important and the project helps companies address this from the start. The EU AI Act, machinery requirements and emerging standards can alter the obligations surrounding an AI application while it is being developed. AIRISE's conformity assessment tool screens against relevant rules and standards, identifies where action may be necessary and produces a report to guide the company. The tool is updated as the regulatory landscape changes, while AIRISE also participates in the standardisation discussions shaping future requirements.

"We do not pretend that our diagnostic tool makes us a lawyer or cybersecurity supplier," adds Thombansen. "It identifies the issue early enough for the manufacturer to act and, where specialist execution is needed, connects the company to expertise capable of completing the work."

The AIRISE ecosystem

That handover reveals why the AIRISE ecosystem is such a vital component of this AI journey for manufacturing SMEs. Research organisations can assess a need, create a demonstrator and develop capability, but an SME may also require sensors, computing hardware, systems integration, finance, training, cybersecurity advice, legal support or a commercial solution already available from the market.

AIRISE brings together universities and research organisations, regional ambassadors and "AI-aware" commercial providers that understand manufacturing.

"We see ourselves as facilitators and analysers, not consultants who have one talk and then run away," Thombansen explains. "The purpose is to accompany an SME from 'I have a problem' through 'Could AI be good for me?' to a small demonstration – and then answer the questions that determine whether progress continues: 'Where do I go now? Where do I try, and who do I trust?'

"European funding made it possible to build that pathway," he continues. "It financed the tools, experiments, training and relationships needed to prove the approach in the real world.

"And that's why we are happy to have had this support from the European Commission," says Thombansen, referring to AIRISE's Horizon Europe funding. "It allowed us to develop these tools without having a paywall right after the first step.

"The result is more than a collection of project deliverables. It is a developed support structure capable of bringing the right people, skills and solutions together around the circumstances of an individual European manufacturing SME."

A lasting mission

As AIRISE approaches the end of its Horizon Europe funding period, the consortium is committed to continuing its work through an association carrying the AIRISE name.

The initial model is deliberately viable at modest scale: paid orientation and consultation services to work with manufacturers, analyse their position and point them towards the right next move.

"We have made it a minimum viable product," Thombansen says. "Investment or additional funding would increase the scale of the operation, but our continuity does not depend on waiting for the next grant."

That transition is the clearest measure of mission accomplished. AIRISE has shown that AI can improve quality, utilisation, planning, resource efficiency and decision-making in very different factories. More importantly, it has created a repeatable way for SMEs to discover what is possible, build the foundations, test value safely, understand the human and regulatory consequences, and find trusted partners for the journey beyond the first experiment.

For Bikas, the motivation is straightforward: "AI is a very useful tool for manufacturing problems when it is done properly."

And for Thombansen, the future direction was clear from the start: "We never spoke about being a project," he says. "We talk about AIRISE as a mission. The EU funding sparked that mission, and the emerging association and the ecosystem we have created will carry it forward."



PROJECT INFORMATION

Project Title
AIRISE – Artificial Intelligence in Manufacturing for Sustainable Applications at SMEs

Project Objective
AIRISE supports European SMEs in the uptake of Artificial Intelligence applied to manufacturing, with the objective of reducing waste and carbon footprint while ensuring the resilience.

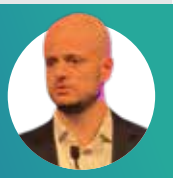
Project Duration and Timing
01 January 2023 – 31 December 2026

Project Funding
Horizon Europe CL4 7.738.776,25 euros

Project Partners
Panepistimio Patron/LMS
Fraunhofer Gesellschaft Zur Forderung Der Angewandten Forschung Ev
Technische Universiteit Eindhoven
European Federation For Welding Joining And Cutting
Asociacion De Investigacion Metalurgica Del Noroeste
Fundacion Tecnalia Research & Innovation
Commissariat A L Energie Atomique Et Aux Energies Alternatives
Lortek S Coop
SINTEF As
Deutsches Forschungszentrum Fur Kunstliche Intelligenz Gmbh
Ceske Vysoke Uceni Technicke V Praze
Viesoji Istaiga Intechcentras
Kliynteh Bulgaria Fondatsiya
Pannon Gazdasagi Halozat Egyesulet

MAIN CONTACT

Name
Dr. Harry Bikas



Phone
+30 2610 910 160

Email
bikas@lms.mech.upatras.gr

Web address
www.airise.eu

AIRISE.EU

This project has received funding from the European Union's Horizon Europe research and innovation program under the grant agreement **No 101092312**.

SMARTER, MORE FLEXIBLE FACTORIES

SoliDAIR: Closing the loop on industrial AI

SoliDAIR is combining visual AI, production data and robotics to make automated quality control more adaptable, trustworthy and practical, helping manufacturers move from promising demonstrations to dependable factory-floor systems. **Andreas Frommknecht**, who leads the Image Processing team at Fraunhofer IPA and coordinates the project, explains how this approach is being developed and tested across four industrial use cases, and why addressing challenges around data, integration and operator trust is essential to wider AI adoption in manufacturing.

While AI has been successfully demonstrated across many industrial settings, the harder question for manufacturers once testing is complete is what happens next. How can a successful experiment become a reliable part of a production line, operate at cycle time, fit in with established factory-floor processes and earn the confidence of the people who will use it every day?

It is this gap between what the tests promise and routine deployment that is at the heart of SoliDAIR – Solid, Rapid and Efficient Adoption of Data, AI and Robotics Application in Production. Bringing together ten partners, the project is developing adaptable technologies for automated quality control and process improvement across three connected fields: visual AI, data and AI, and robotics and AI. The wider goal is to make those technologies replicable beyond the original demonstrations, strengthening the competitiveness and sustainability of European manufacturing.

For project coordinator Andreas Frommknecht of Fraunhofer IPA, the need for this approach is a practical one. AI may have become more familiar during the life of the project, not least through the rapid public uptake of generative tools like ChatGPT and Copilot, but acceptance on the factory floor cannot be assumed. "AI is much more accepted now than it was at the beginning of the project," he says. "But user



Typical manufacturing scene

acceptance is still critical, especially in quality control, where people are used to relying on rule-based systems and still have doubts about AI.''

Those doubts sit alongside several technical and organisational barriers many companies face. Manufacturers may lack established validation criteria or the representative data essential to an AI project. AI can also appear to be a black box whose decisions are difficult to scrutinise and understand. In a mature production environment, moreover, good performance creates a paradox in that there may be plenty of production data, but very few examples of failure.

''Every company wants to produce good-quality parts and avoid bad parts,'' Frommknecht explains. ''But for a supervised AI solution, we also need the bad parts, ideally in a useful balance with the good ones. That can be a very tricky thing.''

A transferable approach

SoliDAIR works across three functions common across manufacturing. Visual AI automates inspections that may previously have depended on the human eye, using cameras and other imaging sensors to identify defects. The data and AI theme works with signals gathered during production, looking for the relationships between process conditions and final quality. Robotics and AI add flexibility to data acquisition and inspection, for example by moving a camera around a complex component to capture the surfaces, threads and angles that a fixed system cannot see.

Physical data collection is therefore indispensable, even though it is not the project's central innovation. ''The core activity is on the AI side and not the data acquisition itself,'' Frommknecht says.

The distinction matters because SoliDAIR is not intended to produce four isolated point solutions. Its industrial cases provide proving grounds for modules designed around common problems such as data scarcity, opacity and difficult inspection planning. ''The aim is not to have specialised solutions only for the use cases addressed in SoliDAIR,'' adds Frommknecht. ''The use cases should show the applicability of use-case-agnostic models.

''And by use-case-agnostic, we do not mean plug and play,'' he continues. ''A model transferred to another component, process or factory will still need to be fine-tuned with relevant data. The value lies in starting with reusable building blocks and a tested method, rather than beginning again from a blank page. The same principle applies to robotic path planning: a module developed to move a camera around one die-cast part can be adapted to plan inspection of a different geometry.''

Learning when defects are rare

While data scarcity can be a challenge shared across sectors,



Optical data acquisition

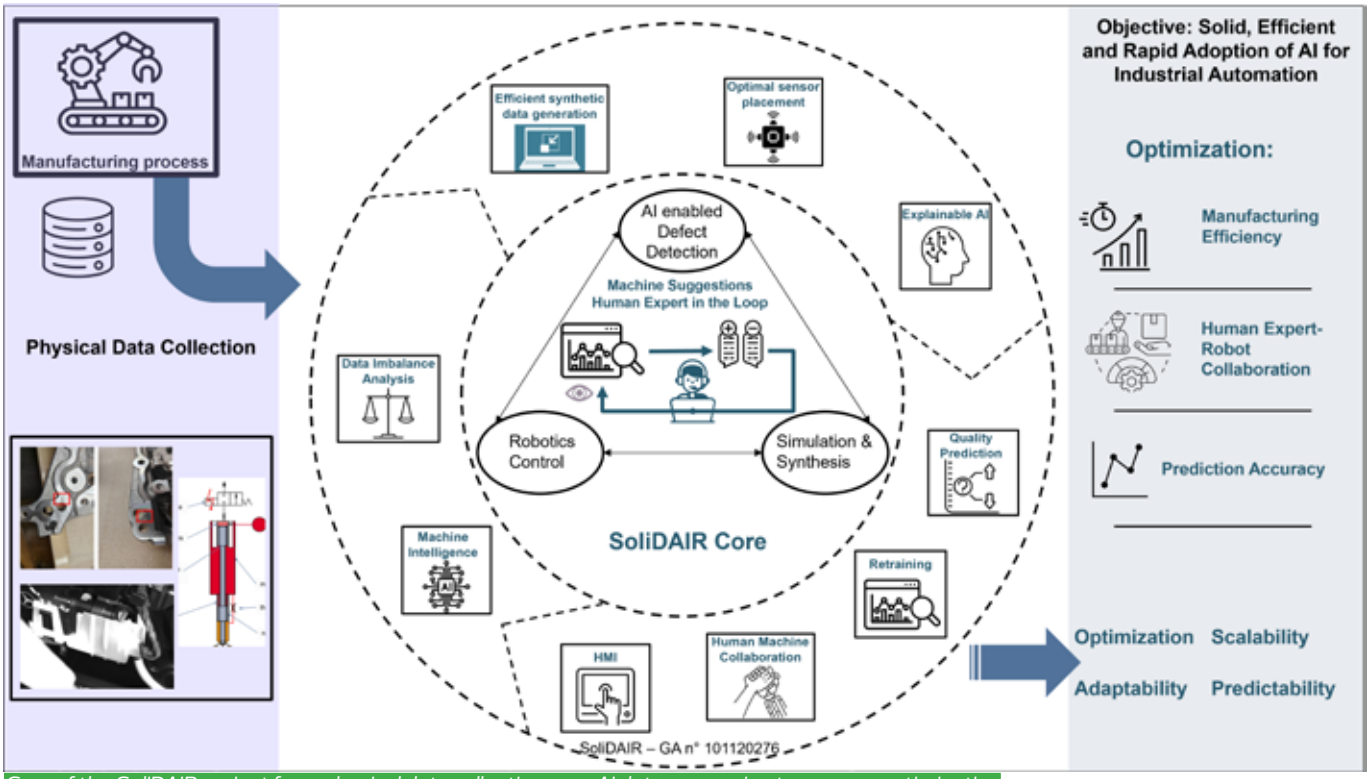
synthetic data can help address this problem by augmenting a limited collection of real images, either by using a model of the component or generating realistic variations from a small set of examples. Anomaly detection can learn the characteristics of acceptable production and flag departures, reducing the need for extensive catalogues of every possible defect.

For sensor-rich mass production, the issue may not be a lack of data overall, but an extreme imbalance within it. When faults occur at parts-per-million rates, the relevant failure signals are vanishingly rare. SoliDAIR combines real and simulated information in hybrid models, reduces the complexity of large sensor datasets and examines where sensors should be placed so that the system captures enough meaningful information without unnecessary instrumentation.

''The tricky part is really to figure out which data is missing and how we can generate it, at best synthetically,'' Frommknecht says. ''Once suitable data is available, many established AI architectures can deliver the required analysis; the greater industrial challenge is assembling evidence that genuinely represents the process and its edge cases.''

Building trust into the system requires more than accuracy. In the visual AI theme, explainability tools such as heat maps show which area of an image influenced a model's decision. That helps an operator see whether the system has focused on the suspected defect or learned for example an irrelevant shortcut from the background. ''The human has the possibility, for every decision taken by the AI, to take a closer look,'' he says, ''and to understand whether the focus was on the right part of the picture and not on a completely false one.

''Confidence scores add another layer by indicating how certain the model is, helping to identify cases that warrant human review. Together with clear interfaces, continual evaluation and retraining, these measures turn



Core of the SoliDAIR project from physical data collection over AI data processing to process optimisation

trustworthiness from a nice-to-have into a set of practical design choices."

From simulation to the production line

Robotic inspection introduces a different form of uncertainty: the system must act in a particular physical space, often with very specific requirements and challenges. SoliDAIR uses simulation and digital twins to plan camera positions and robot paths offline. Without this digital test environment, engineers would have to mount the equipment and repeatedly move it around a real component, testing whether each position provides the necessary view. It also helps engineers manage collision risks and avoid production losses while tests take place.

"By using the simulation module, you can know which positions will work and which will not," Frommknecht explains. "It is much faster than doing it in the physical world. When you are ready, you deploy it to the physical system." This simulation-to-real approach accelerates commissioning and reduces risk before the robot reaches the line.

Validation

All these technologies have been validated in four realistic industrial scenarios. An automotive door-module case by BROSE concentrates on stationary visual inspection and synthetic data. An aluminium high-pressure die-casting case by CIE combines visual AI with robot-mounted cameras to examine complex surfaces and features. A high-volume precision-component case by BOSCH tackles the difficulty of finding quality-relevant signals among very large datasets when failures are exceptionally rare. A multi-step gearbox-assembly case by AUTFORCE uses AI to interpret and analyse audio signals from the assembly process.

Together, the four cases expose practical factory-floor issues that laboratory benchmarks alone cannot capture. A false positive can divert a good part for review, but a false negative allows a defective part to continue. "In applications where escaped defects carry the greater risk, models can be tuned to prioritise reducing false negatives, accepting additional false alarms where appropriate. Frommknecht says. "Models can therefore be tuned to reduce escaped defects, accepting more false alarms where necessary.

"No inspection system, however, whether human, rule-based or AI, is entirely free of ambiguous boundary cases," continues Frommknecht. "A defect threshold for example of one millimetre will inevitably have examples just below and just above the limit. The objective is to manage that uncertainty transparently while meeting production requirements."

The real challenge is balancing reliable defect detection with the need to maintain production speed. "You do not want quality control to decrease the speed of the production process," he adds. "The quality control has to run within the cycle time."

Commercialisation, equipment costs and long-term reliability are also being considered with the industrial partners, although the research priority has been to establish what

works with a reasonable implementation effort. That distinction keeps the results grounded without presenting every module as an off-the-shelf product.

Even at this pre-commercial stage, the achievement is tangible. Across four markedly different production environments, SoliDAIR has shown how a common family of AI, data and robotics technologies can be adapted to address distinct quality-control challenges, while accounting for practical requirements including cycle time, false alarms, explainability and operator review. The result is a tested route from isolated demonstrations towards more repeatable and dependable factory-floor deployment.

Implementation is a human process

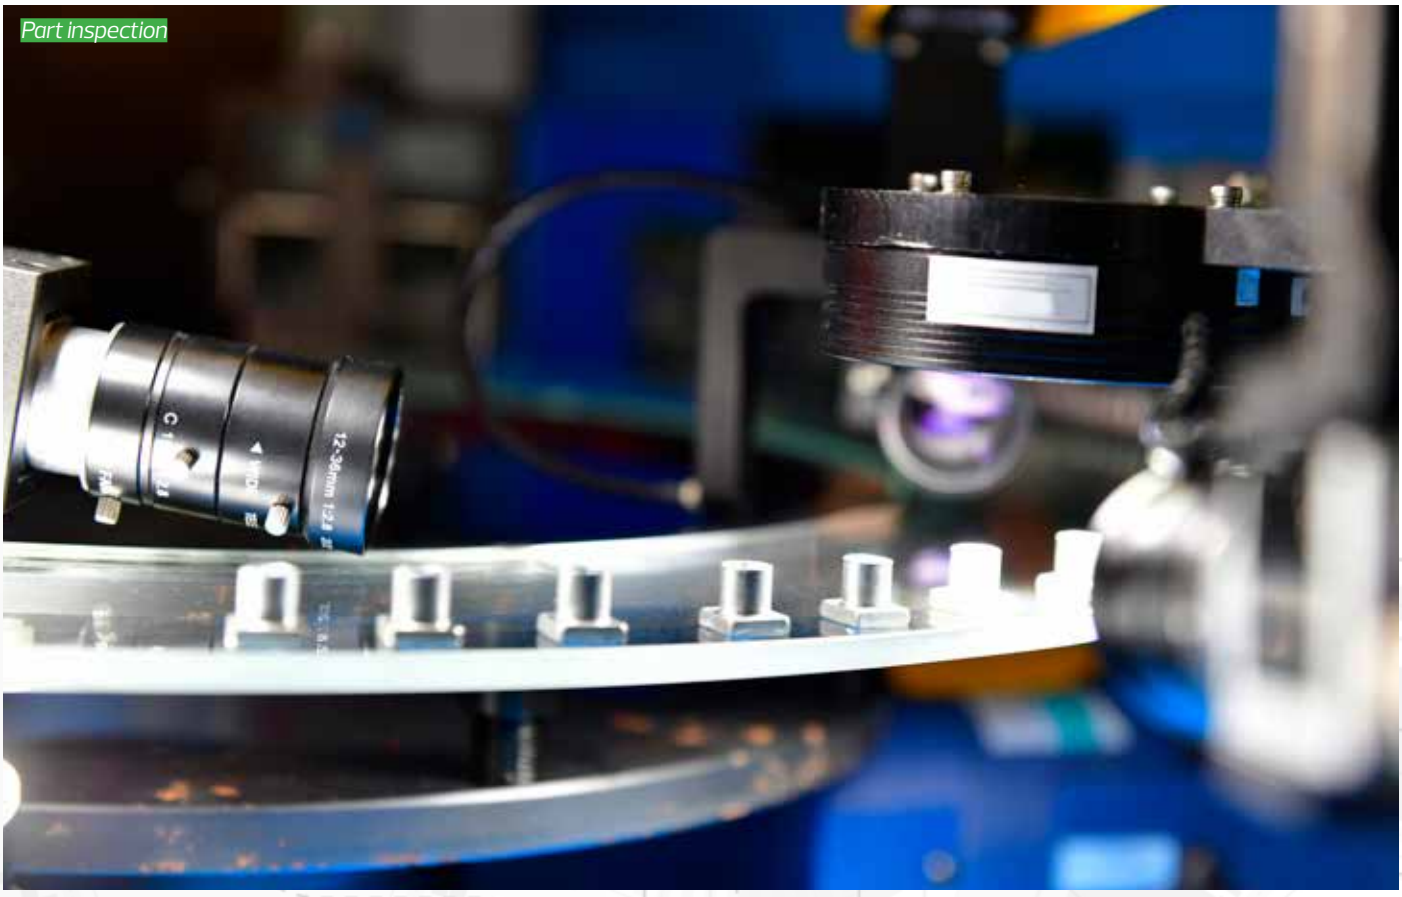
SoliDAIR's trials reinforced an important lesson shared across all three technology themes: deployment is not complete when a model is installed – explainability and flexibility are key. Systems must be straightforward to use and continually evaluated and refined as products, materials and operating conditions change. Modularity makes this manageable, while retraining the models keeps performance firmly in line with production realities.

Operator involvement is equally important and SoliDAIR has produced training and upskilling materials while also considering human-machine interfaces alongside model performance. "It is very important to include the humans, take them on board and present AI as a help, not as a competitor," Frommknecht says. "In practical terms, that means allowing operators to move away from repetitive inspection and concentrate on more complex tasks where human judgement adds greater value."

This demands collaboration between process engineers, production specialists, data scientists, AI developers, robotics teams and operators. Large manufacturers may be able to assemble that expertise internally; smaller businesses, however, are more likely to need external support. SoliDAIR does not prescribe a provider but brings the required capabilities into the conversation from the start.

"We have developed a step-by-step guide that takes manufacturers from scoping a problem through data acquisition and augmentation, architecture selection, model training, testing and validation, explainability, improvement and human-machine integration," Frommknecht explains. "It is organised around visual AI, data and AI, and robotics and AI, and has links to relevant datasets, code, pre-trained models for companies to use, while also providing sources for further information and expertise."

The guide is more than a catalogue of SoliDAIR technology. Manufacturers are able to choose an open-source model or another supplier's system; the aim is to clarify the path from an identified quality challenge to a working deployment. "It



Part inspection

is not a simple readiness score; It gives you the way." Frommknecht points out. "The users still need to check if they are able to acquire the right data, have the fitting expertise and the necessary budget. So, the usage of the step-by-step guide makes the route clearer: You know which people you need, you can check what you have in your company, and if necessary, you know which expertise you need to search for on the market.

"By showing what expertise is needed at each stage, the guide also allows manufacturers to involve AI specialists more selectively, helping to control costs and lower the entry barrier for smaller companies," he adds.

Wider adoption

The project's tangible legacy combines adaptable modules, free datasets, documentation, software, training resources and the implementation guide, with much of the material available through the project website, the Zenodo repository and GitHub repository. The four use cases demonstrate that the components can address distinct manufacturing environments while preserving a common methodological core. Transfer will always require fine-tuning, but the resources provide a clearer starting point and reduce the uncertainty surrounding early implementation.

This all matters because, in Frommknecht's view, industrial AI is moving from an option to a competitive necessity. "Europe may trail international competitors in some prominent areas of AI, but we have deep manufacturing knowledge and valuable production data. Connecting those assets with trusted AI is an opportunity to build on this existing industrial strength.", he says.

Realising that opportunity, Frommknecht argues, will require sustained investment from both the private and public sectors if Europe is to turn its manufacturing expertise and production data into a lasting AI advantage.

While not directly part of SoliDAIR, its work could provide a springboard towards embodied AI: systems able not only to identify drift and recommend revised parameters, but potentially to adjust a manufacturing process online, subject to appropriate safeguards and human authority.

"The next big step is not only to check quality; it is the production of quality," Frommknecht says. "That means bringing the quality information back into the control loop of the machine." Because such systems act in the physical world, the risks are greater. "A wrong action can produce defective parts or damage equipment," he adds. "Human oversight therefore remains fundamental – the human should be the final decision-maker."

Following the conclusion of SoliDAIR, downloads of its resources and participation in its webinar suggest that manufacturers are seeking precisely this kind of practical bridge. Its four validated demonstrators and package of reusable resources already provide a tangible result; the longer-term measure of success will be whether manufacturers use them to move beyond isolated demonstrations – building AI systems that fit their production lines, explain their judgements and turn Europe's manufacturing knowledge into sustained advantage.

PROJECT INFORMATION

Project Title
Acronym: SoliDAIR, full title: Solid, rapid and efficient adoption of Data, AI & Robotics applications in production

Project Objective
SoliDAIR brings together 10 manufacturing, technological and research experts to accelerate the uptake of AI and Robotics in European manufacturing, while using Data as an enabler, by co-developing and demonstrating tailored solutions to digitalise and automate visual inspection and physical testing, enable predictive quality control and process optimization.

Project Duration and Timing
1 October 2023 to 30 September 2026

Project Funding
HORIZON-CL4-2022-DIGITAL-EMERGING-02, Topic: HORIZON-CL4-2022-DIGITAL-EMERGING-02-05, Total eligible costs: 3 581 121.25, max grant amount: 2 988 331.76

Project Partners
Fraunhofer Gesellschaft zur Forderung der Angewandten Forschung e.V., Brose Fahrzeugteile SE & Co. Kommanditgesellschaft, Fundacion CIE I+D+I, Bosch Sanayi Ve Ticaret A.S., Autforce Automations GmbH, Siemens Industry Software NV, UG Systems GmbH & Co. KG, Tech Hive Labs Astiki Mi Kerdoskopiki Etaireia, Virtual Vehicle Research GmbH, I2m Unternehmensentwicklung GmbH

MAIN CONTACT

Name
Andreas Frommknecht



Phone
+49 711 970-1818

Email
andreas.frommknecht@ipa.fraunhofer.de

Web address
https://www.solidair-project.eu/

SoliDAIR

This project has received funding from the European Union's Horizon Europe research and innovation program under the grant agreement **No 101120276**.

ADVANCED SENSORS AND PHOTONICS

LASTSTEP: Bringing mid-infrared sensing onto silicon

Gases and liquids leave distinctive optical fingerprints in the mid-infrared, but the components needed to read them remain relatively expensive and difficult to integrate. The €4.18 million Horizon Europe LASTSTEP project is developing a silicon-compatible platform that brings light sources, detectors and photonic circuits together on 200 mm wafers. **Vincent Reboud** of CEA-Leti explains how laboratory advances are being shaped into useful devices that could eventually be manufactured at scale for environmental monitoring, food analysis, health and industrial safety.

Light can identify a substance without touching it. In the mid-infrared region of the spectrum, molecules absorb particular wavelengths according to their chemical composition. These patterns act like fingerprints: choose the right wavelength and a sensor can look for methane in the air, distinguish an authentic olive oil from a tampered product or, in time, detect markers connected with health and biological processes.

This principle is established, and specialist mid-infrared sources and detectors are already available. The constraint with it, however, is the way those components are made. High-performing devices often rely on compound semiconductor materials that do not fit naturally into the silicon manufacturing infrastructure used for microelectronics and much of modern photonics. They can be attached to a silicon platform, but this heterogeneous approach adds fabrication, alignment, and packaging steps. Cost and complexity consequently limit how widely mid-infrared sensing can be deployed.

LASTSTEP is pursuing a different route by using silicon-germanium-tin alloys, known as SiGeSn, to develop an all-Group-IV photonic platform capable of emitting, guiding and detecting light between 2 and 5 micrometres. Its particular focus is around 3.3 micrometres, a region relevant to methane and oil sensing. Because silicon, germanium and tin are all Group IV elements, the platform can be designed around CMOS-compatible processes: the established tools and production controls used to make silicon chips. This avoids joining a separate compound-semiconductor source and offers a route towards a fully integrated system made in a silicon foundry.

"We are trying to build something that does not exist: a mid-infrared sensing platform that could, in the future, be made in a foundry," says Vincent Reboud of CEA-Leti, the project's coordinating organisation. "Each molecule has a very specific resonance. If you have the right light, you will be able to detect methane, carbon dioxide or some biological molecules."

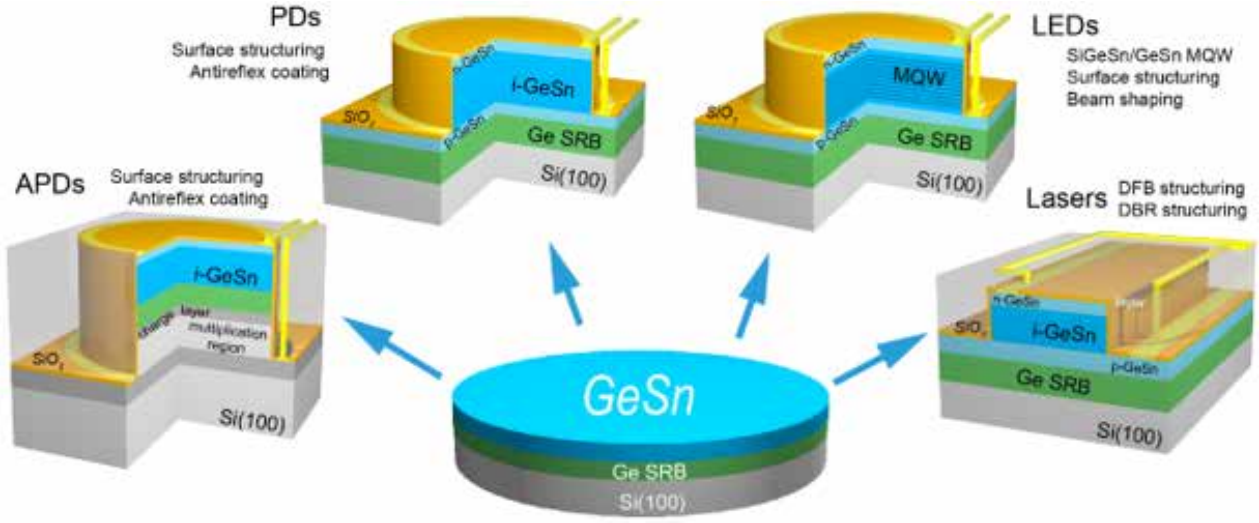


Figure 1: PDs, APDs, LEDs and lasers from one GeSn stack on Si(100)

The addition of tin changes the optical behaviour of germanium. With the appropriate composition and structure, the alloy can acquire a direct bandgap and emit light efficiently at mid-infrared wavelengths. Varying the tin content also provides a means of tuning the wavelength for different targets. LASTSTEP's task is to turn that material opportunity into a practical platform containing light-emitting diodes (LEDs), lasers, photodetectors and avalanche photodetectors, first as discrete components (Figure 1) and then as devices connected through photonic integrated circuits (PICs).

LASTSTEP is testing the platform through two application-led use cases. ELICHENS is using the components to measure methane in gas, while DIAFIR is investigating whether mid-infrared light can authenticate premium olive oil and detect adulteration. Each use case is being developed first with separate components and then as an integrated photonic link, providing a practical route from individual devices to a compact sensing system.

A missing piece

Silicon photonics is already established in communications and data centres, where photonic circuits carry and manipulate light at telecom wavelengths. LASTSTEP is seeking to extend the manufacturing logic of that industry into the mid-infrared. Its aim is a reusable technology from which many sensors could be designed.

The project's case therefore rests on manufacturability as much as absolute device performance. Specialist compound-semiconductor sources may retain advantages where an application needs the greatest power or sensitivity. GeSn's value will come from reaching a level of performance that is sufficient for useful applications while offering a credible route to integration and volume production.

"What we aim for with germanium-tin is maybe not the very highest optical performance, but decent performance that can be produced in very large numbers and put into different sensors," Reboud explains. "The gain will be more on the cost, and we could make the use of mid-infrared sensors much more widespread."

It is that proposition that brings a fundamental materials challenge for the project to tackle. "Growing GeSn directly on silicon places together crystals with different lattice dimensions," Reboud continues. "The mismatch creates defects, or dislocations, which can reduce the output of a light source and increase noise in a detector. We have to control those effects while developing contacts, device structures, waveguide coupling and packaging that work as a complete system."

The manufacturing scale matters, and LASTSTEP is developing processes for 200 mm, or eight-inch, wafers (Figure 2). Existing specialist mid-infrared technologies are generally produced on much smaller substrates, whereas 200 mm processing can

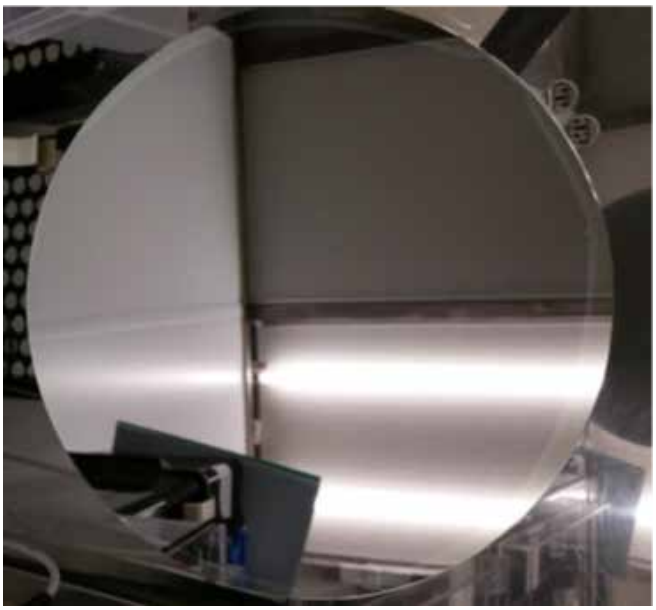


Figure 2 : Example of GeSn thick layer with 16% Sn content epitaxied on 200 mm Si wafer.

use highly repeatable equipment and established production disciplines. More devices can be made in each run, and the process has a clearer path towards foundry transfer.

Reboud describes this as applying the microelectronics vision to a new spectral range. Large-wafer fabs use controlled tools continuously, building consistency and reducing unit cost. "This could be game-changing, because there is no comparable III-V or II-VI technology on eight-inch wafers," he says. "It allows very good reliability and repeatability and can massively reduce the cost of fabrication."

Closing the development loop

LASTSTEP brings together materials growers, device physicists, modelling teams, cleanroom specialists, a photonic design and test house, and two sensor SMEs. Their work cannot proceed as a simple sequence in which one group completes its task and passes the result to the next. GeSn remains a young semiconductor material, with important physical parameters still being refined. Design, fabrication and characterisation must continually inform one another across all involved stakeholders.

The modelling teams can predict the behaviour of LEDs, lasers and photodetectors, but the more fundamental description of the material is still developing. Reboud says the consortium is not yet fully predictive at that level. Instead, the partners grow the material, fabricate devices, measure their behaviour and return the results to the modellers. The models are adjusted, and the next material stack or device design can be better targeted.

"We do epitaxy, we do fabrication, we do characterisation, and then we go to the modelling people," he says. "We have this loop to understand and evaluate the models against the

characterisation. There is a very strong interaction between modelling, characterisation and fabrication.''

Those application partners add a second development loop, eLichens for its methane sensing, and DIAFIR for its sensing work for olive-oil authentication. They define what the sources and detectors must achieve, but they also expose requirements that do not appear in a device-performance table. A component must have the right dimensions, package and optical alignment to fit a practical sensor. Its output must be usable by the rest of the system.

This has shaped the project's two-stage approach. Discrete devices allow the sensor companies to assemble test benches and compare GeSn components with established products (Figure 1). Integration is the next, more ambitious step: placing a source and detector on a photonic platform and coupling the light into waveguides that carry it through the chip (Figure 3). The result should be smaller, reduce alignment demands and simplify packaging, while making the platform more adaptable to different sensing tasks.

Reboud is candid about the difficulty. ''Integration will bring a lot of versatility to spread the technology, but there is a big, big challenge,'' he says. At the time of the interview, the team had coupled light into waveguides and fabricated samples intended to test whether a light source and photodetector could operate together on the same photonic chip. Characterisation of those samples was under way.

Evidence from working devices

The platform contains four principal device types: LEDs, conventional photodetectors, avalanche photodetectors and lasers. The first three operate at room temperature. The hardest objective is an electrically pumped laser that can do the same.

LASTSTEP has already demonstrated room-temperature lasing at 3.3 micrometres under optical pumping. This is an important proof that the material can provide optical gain at the target wavelength, but optical pumping requires another light source and is not the practical end point. The consortium has also demonstrated continuous-wave electrically pumped GeSn lasers at low temperature. Raising their operating temperature remains a longer-term challenge, while the immediate project effort is concentrated on integration, packaging and end-user evaluation.

Progress in the less spectacular components may be more directly relevant to the first sensors. Reboud reports that the optical power of the LEDs has increased by a factor of seven since the start of the project, taking them to the application level requested by the partners. New photodetectors are also close to the specified detectivity, a measure that combines their response to light with the noise they produce.

''All the discrete components are very close to, or at, the performance required by our end users,'' he says. ''We have also made the first demonstration of a gas sensor with germanium-tin components. That shows these devices can be used for a practical application.''

The methane demonstrator makes the remaining gap tangible, with initial trials producing an estimated detection limit of around 7,000 parts per million, compared with a target of approximately 5,000 ppm for the intended practical use. Figure 4 presents how the sensor components match the methane absorption. Methane absorbs most strongly around 3.3 µm,

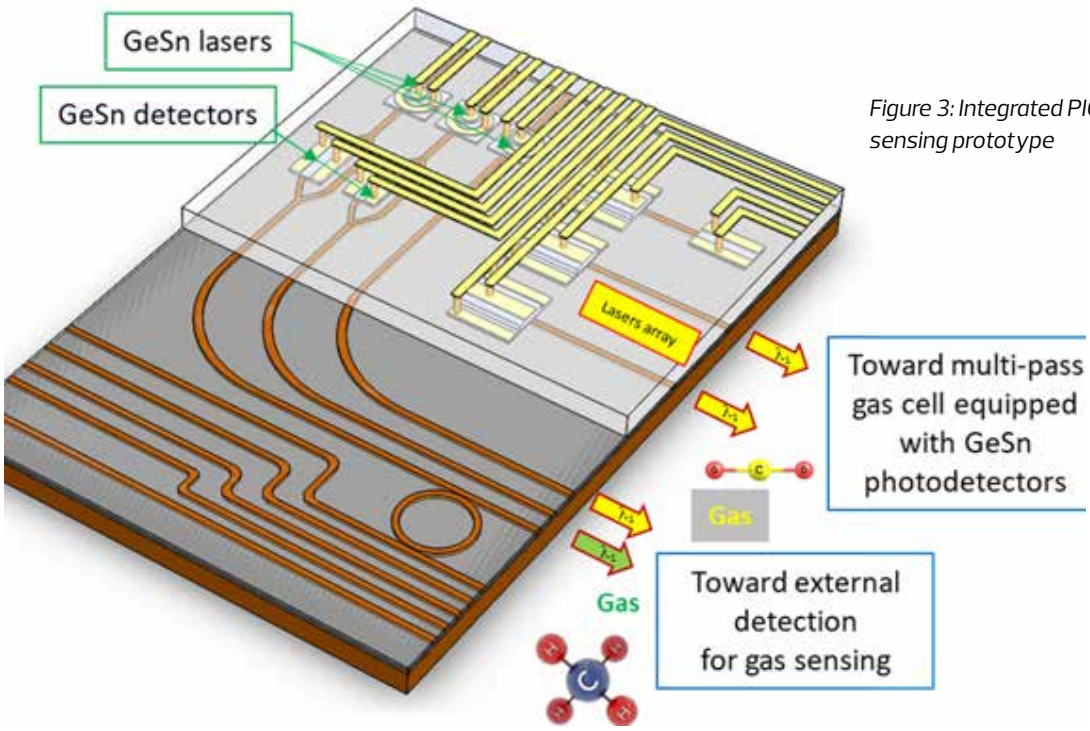


Figure 3: Integrated PIC gas-sensing prototype

which is also where the GeSn LED with the right Sn content amount emits the most light. A filter lets through only this region (about 3.2–3.45 µm) and blocks the wavelengths around. The GeSn photodiode can detect light up to about 3.6 µm to cover the whole filter range. The methane-selective optical filter slightly reduced the signal, while detector noise and LED output constrained the result. New generations of both devices were delivered for further tests during the summer, and Reboud expected their higher output and detectivity to close much of the difference.

The intended product is deliberately a relatively low-cost methane detector rather than the most sensitive instrument available. An LED offers less spectral precision than a laser, but it may suit applications where affordability, compactness and deployment at many points matter more than detecting the smallest possible concentration.

For olive oil, light from the source is coupled into a thin fibre. A small part of the guided field interacts with the liquid around it, creating a spectrum that can be analysed for signs of authenticity or tampering. Results from those trials were not yet available at the time of the interview. The distinction is important: LASTSTEP has produced strong component evidence and an initial gas-sensing proof of concept, but the complete integrated platform and both end-use demonstrations still require final validation.

Changing the science

Not all the project's challenges originated in the laboratory. GeSn growth had relied on digermane, a specialist gas precursor supplied by a single producer. When that supplier discontinued production during the project, the consortium faced a potential interruption to its material programme.

The two institutes responsible for GeSn growth responded by working together on a process based on germane, a mainstream precursor used in microelectronics. Changing the source gas alters the growth chemistry and meant rebuilding process knowledge, but the partners achieved material performance comparable with the previous route. Further optimisation is needed for growth rate and reproducibility, yet the unplanned change may ultimately strengthen the manufacturing case by removing dependence on a scarce input.

"There was only one provider doing digermane, and they stopped production during the project," Reboud recalls. "The institutes really came together and developed a new epitaxy process with germane. We obtained very similar performance, and germane is a mainstream precursor for microelectronics."

The project has also learned when to stop pursuing one line of research and redirect effort towards a demonstrator. User requirements evolved once complete sensor benches could be tested, and the consortium adjusted priorities between light sources, detectors and packaging. Reboud says that, at one stage, some developments were paused so that more resources could be placed behind the end users' demonstrators. The deciding measure became whether the partners could build and test the system the application required.

This discipline has helped the consortium translate between research and use. A laboratory may be able to fabricate a high-performing device, but an SME needs it in a form that can be mounted, aligned and replaced. Several practical boundaries only emerged through repeated discussion about what each partner could supply and what the next partner could realistically handle.

A route through the foundry

LASTSTEP is establishing an experimental platform rather than a production-ready commercial service. Moving from the current prototypes to industrial supply will require technology transfer into a foundry, followed by further work on process repeatability, yield and long-term reliability. Those activities sit beyond the project's exploratory scope, and they will need sustained investment and an industrial actor willing to adopt the technology.

Reboud nevertheless sees no fundamental roadblock, as the necessary categories of growth, lithography, etching and cleanroom equipment already exist. The challenge, he believes, is to qualify a GeSn process around them and reproduce the results outside the research institutes that developed it. "Transferring to industry is a big job, it will cost a lot of money, and you need actors interested in it," he says. "But the tools exist in the cleanroom. We want to show the capability of the material, provide the base and then build the transfer to industry.

"If that transfer succeeds, a foundry could hold a qualified material stack and process flow that customers use as a platform," he continues. "A large company could design a specialised sensor without creating the semiconductor technology from the beginning. An SME could gain access to integrated mid-infrared components without owning an epitaxy reactor or a silicon line. Different products would select an emission wavelength, optical circuit, package and analysis method for the molecule they need to measure."

Environmental monitoring appears to Reboud to be the clearest early opportunity, with distributed sensors that could monitor greenhouse gases, urban air quality or safety conditions inside buildings and industrial plants. He also sees further application possibilities for food analysis, which could extend from olive oil to other questions of authenticity and quality. Medical possibilities, meanwhile, include breath analysis and other non-invasive tests, although complex biological molecules will demand carefully designed systems. "The same material research has also opened questions in quantum and spin-related physics that were not central to the original application plan," he adds.

The value of the platform is therefore its capacity to support uses the consortium cannot define in advance. "LASTSTEP

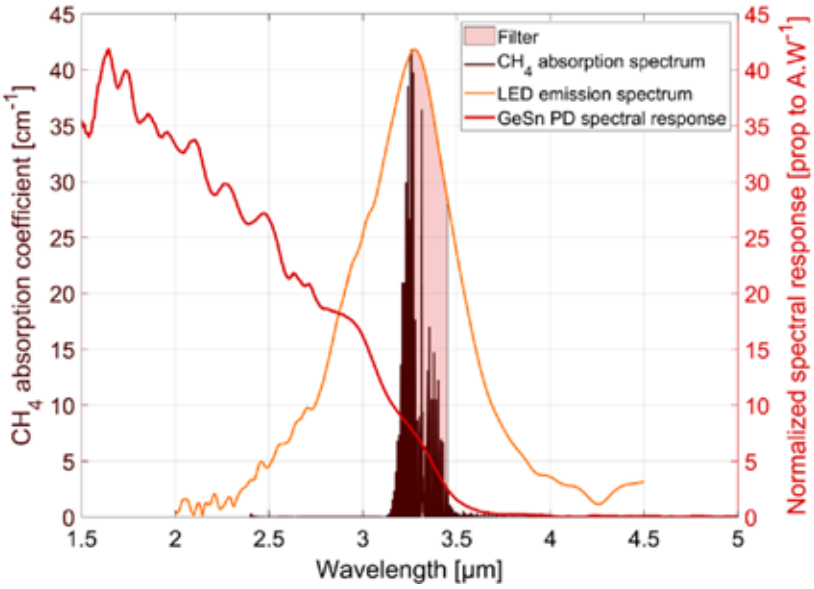


Figure 4 : CH_4 absorption coefficient (left axis) compared with the normalised LED emission, band-pass filter transmission and GeSn photodiode spectral response (right axis), showing their overlap around the 3.3 µm methane band.

is developing the component technology," Reboud explains. "Future companies and researchers will combine the appropriate wavelength, sensing geometry and signal analysis for their own targets. Lower-cost manufacture would make it feasible to place sensors in more locations, turning mid-infrared analysis from specialist instrumentation into a more widely available source of environmental, safety and health information."

For Europe, the opportunity also has industrial value – the continent already has leading GeSn research teams and the cleanroom capability needed to develop the technology. A European foundry process would convert that knowledge into an accessible manufacturing asset rather than leaving each application to assemble a supply chain of specialist components.

The final months of LASTSTEP are focused on the evidence needed to justify that next step: benchmarks from the methane and olive-oil demonstrators, and confirmation that the light source and detector can operate together through the photonic link. This evidence should show whether the central proposition is sound, while later programmes address performance, reliability and commercial demand in greater depth.

"A big success would be to have germanium-tin in a European standard silicon line," Reboud concludes. "It could become a key technology for a mid-infrared photonic platform that is in the foundry and widely used."

It's an ambition that brings new physics and manufacturing pragmatism together, with the project asking whether a material developed by a relatively small research community can supply the missing light sources and detectors for silicon photonics in the mid-infrared. A credible route by which many useful devices can be made would be its most important result.

PROJECT INFORMATION

Project Title
LastStep : group-IV LASer and deTectors on Si-TEchnology Platform

Project Objective
LASTSTEP aims to deliver efficient group-IV light sources monolithically integrable with CMOS electronics, enabling fully silicon-based photonic integrated circuits for the mid-infrared range.

Project Duration and Timing
50 months. Started the 01 Jan 2023

Project Funding
Horizon Europe

Project Partners
CEA (France) – Commissariat à l'Énergie Atomique et aux Énergies Alternatives (project coordinator)

JUELICH (Germany) – Forschungszentrum Jülich GmbH,

POLIMI (Italy) – Politecnico di Milano

C2N (France) – Centre de Nanosciences et de Nanotechnologies, participating as Université Paris-Saclay

UNIVLEEDS (United Kingdom) – University of Leeds

ELICHENS (France) – eLichens

VLC (Spain) – VLC Photonics S.L.

DIAFIR (France) – DIAFIR (Rennes-based SME)

MAIN CONTACT

Name
Vincent Reboud



Phone
+ 33 (0)4 38 78 4167

Email
vincent.reboud@cea.fr

Web address
www.leti-cea.fr/cea-tech/leti

This project has received funding from the European Union's Horizon Europe research and innovation program under the grant agreement **No 101070208**.

ADVANCED SENSORS AND PHOTONICS

SMARTSHAPE: Beyond the blood pressure snapshot mission

SMARTSHAPE is developing a miniature, battery-free implant designed to provide clinically accurate, beat-by-beat blood pressure monitoring for patients in hospital and at home over extended periods. University of Galway Programme Manager **David Connolly**, SMARTSHAPE Commercial Lead **Ciarán Duffy** and Merakoi CEO and Co-founder **Kevin Michels-Kim** explain why delivering this innovation depends not only on sensor accuracy, but on manufacturing discipline, patient trust and a viable route to take it into everyday care.

Hypertension affects an estimated 1.4 billion people worldwide and is linked to around nine million deaths each year, yet it remains poorly controlled in most of those diagnosed. SMARTSHAPE Programme Manager David Connolly believes part of the problem lies in how blood pressure is measured. Although it fluctuates throughout the day and night, treatment decisions still depend heavily on occasional readings that capture only a moment in time.

The conventional consultation usually involves a clinician taking three of these readings, effectively sampling three heartbeats at that moment of that day. An ambulatory blood pressure monitor offers more information but typically records at intervals of 15 to 30 minutes over 24 hours. Connolly likens this to judging a favourite film from one frame, or perhaps a small selection of frames. "You don't really get the full story if you just take a few frames from one single movie," he says.

Continuous measurement is possible through an arterial line, but that is an invasive hospital procedure reserved for settings such as intensive care and surgery. Optical wearables, meanwhile, infer pressure from light passing through the skin, where factors including skin-tone variation and background light can affect consistency. Devices that inflate at the wrist or apply pressure through tonometry can also be uncomfortable and may influence the measurement they are trying to take.

SMARTSHAPE is a project addressing all these issues head-on with a deliberately clinical proposition: "We're not going to approach that market as a wellness device," Connolly explains. "We're going to be a prescribed, reimbursable clinical device designed for people whose hypertension is difficult to control or whose risk justifies closer surveillance, including post-stroke patients and people with chronic kidney or cardiovascular disease.

"SMARTSHAPE is developing a device that will acquire beat-by-beat information outside hospital without requiring an indwelling arterial line," he adds.

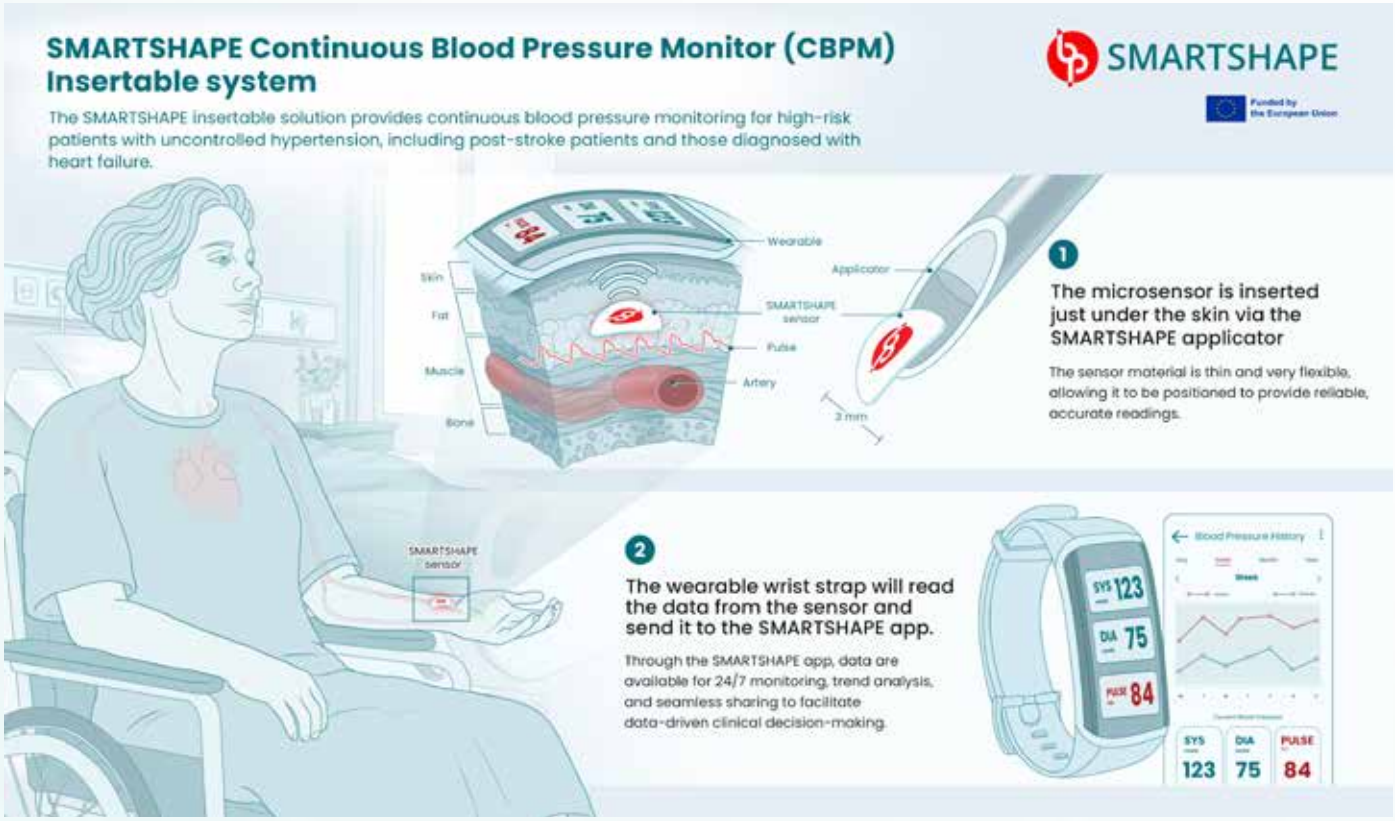
The project was built around a proof-of-concept sensor developed at the University of Galway and already operating in the lab demonstrating that the underlying sensing principle could work. Galway now coordinates a consortium combining technical development with expertise in patient engagement, regulation, sustainability and manufacturing. This breadth reflects the challenge the project has since encountered: a promising sensor is only the starting point for a workable medical system.

Designing for the body

Turning that proof of concept into a clinically useful and manufacturable device exposed the project's first major challenge: repeatability. "At first we had a planar design which was put through a manufacturing process optimised in our University of Galway lab" Connolly explains. "Sensor to sensor, we were seeing huge variability and an effect on the actual sensitivity."

The team used finite-element and in-silico modelling to rework the geometry, moving from an initial proof of concept geometry to a geometry we can now manufacture consistently". Connolly says changes to the component structure "also made the sensor more robust, less susceptible to fluid ingress and capable of a stronger sensor performance"

One of the project's most significant design decisions was not to proceed with the planned integration of a temperature-responsive shape-memory polymer into the blood pressure sensor. The material was intended to allow the sensor to be curled for insertion and then open into a predefined shape at



body temperature. Modelling showed, however, that this coiling and shape shifting would prevent the sensor from retaining consistent sensitivity.

The University of Birmingham, also a project partner, continued to develop the shape-shifting material for other applications, while SMARTSHAPE's pressure device adopted polydimethylsiloxane, or PDMS. "It's a well-known polymer. It's used in implantables already, so it is well characterised, safe, and it works quite well across our blood pressure range of interest," Connolly explains. "Preserving the performance that matters took precedence over preserving an attractive feature from the original concept."

Biology complicated the engineering further when SMARTSHAPE moved from bench testing into preclinical studies. The programme used a sheep model to evaluate the sensor's safety and efficacy, but reproducing its intended position in the human body was not straightforward. The device is designed to sit beneath the skin close to the radial artery at the wrist, for which a sheep has no directly comparable anatomy. As Connolly puts it, "a sheep doesn't really have a human wrist". The team therefore had to develop a new surgical model, establish its feasibility in a pilot study and then progress to acute and chronic studies.

Longevity testing proceeded in parallel, with materials and complete sensors exposed to water, physiological media and cell-culture conditions for up to 24 weeks, with particular attention paid to the polymer's elastic modulus – an important property to manage as it allows the sensor to detect the pulse pressure. Connolly reports that "water had the biggest impact", while exposure to physiological media produced a change of only about five per cent in the material's modulus.

"The resulting change in sensitivity was approximately 10 to 13 per cent, but, crucially," he adds. "It still functioned perfectly in that range of blood pressure we were interested in."

Sterilisation was another important stage of development that could not simply be treated as a downstream production step because each method had the potential to alter the material and its sensing performance. The project examined autoclaving, hydrogen peroxide, plasma treatment and gamma irradiation. "We landed on gamma radiation as the one that caused the least amount of change," Connolly explains, whereas autoclaving proved "completely unsuitable" for the material.

That system must also function in a body that rarely remains still. Movement can disrupt the relationship between the implanted sensor and its external reader, so SMARTSHAPE is combining improvements to the physical design of the wearable with additional sensing that can recognise movement and flag periods when a reliable reading cannot be taken. The team has characterised different wrist movements and body positions, including walking, standing, sitting and sleeping. The objective, Connolly says, is to locate "good long runs of accurate, non-impacted data within the continuous record".

"Movement will always be an issue," he acknowledges, "but the

volume of beat-by-beat information changes the calculation. A conventional ambulatory monitor may take a reading only once every 15 to 30 minutes." Connolly calculates that SMARTSHAPE could collect more readings than a conventional 24-hour ambulatory monitor within approximately three minutes.

The clinical signal

The implant itself is in the order of size that will allow the sensor to be implanted using a minimally invasive delivery device or applicator just under the skin above the radial artery in the wrist. "Insertion of the device could ultimately be done in a GP clinic or by an advanced nurse practitioner," says Connolly. "The sensor has no battery, and an external reader worn over the implantation site energises the sensor while also receivings the response signal that is converted to pressure values in mmHg.

"The sensor is a dynamic and passive sensor converting the deformation of the sensor to signal that is detectable by the linked wrist worn reader system continuously," he continues. "When the sensor is exposed to the mechanical force of the nearby artery's pulse wave it undergoes a cycle of varying deformation which is directly linked to the strength of the force, or in other words the patient's blood pressure.

"During calibration in the lab, the device is exposed to known pressures so that the signal shifts can be mapped to pressure values. Signal processing then filters noise and converts the response into a continuous pulse wave rather than only a systolic and diastolic number. It's not just two static values," he concludes. "You get the continuous pulse wave being read with the electronic reader system."

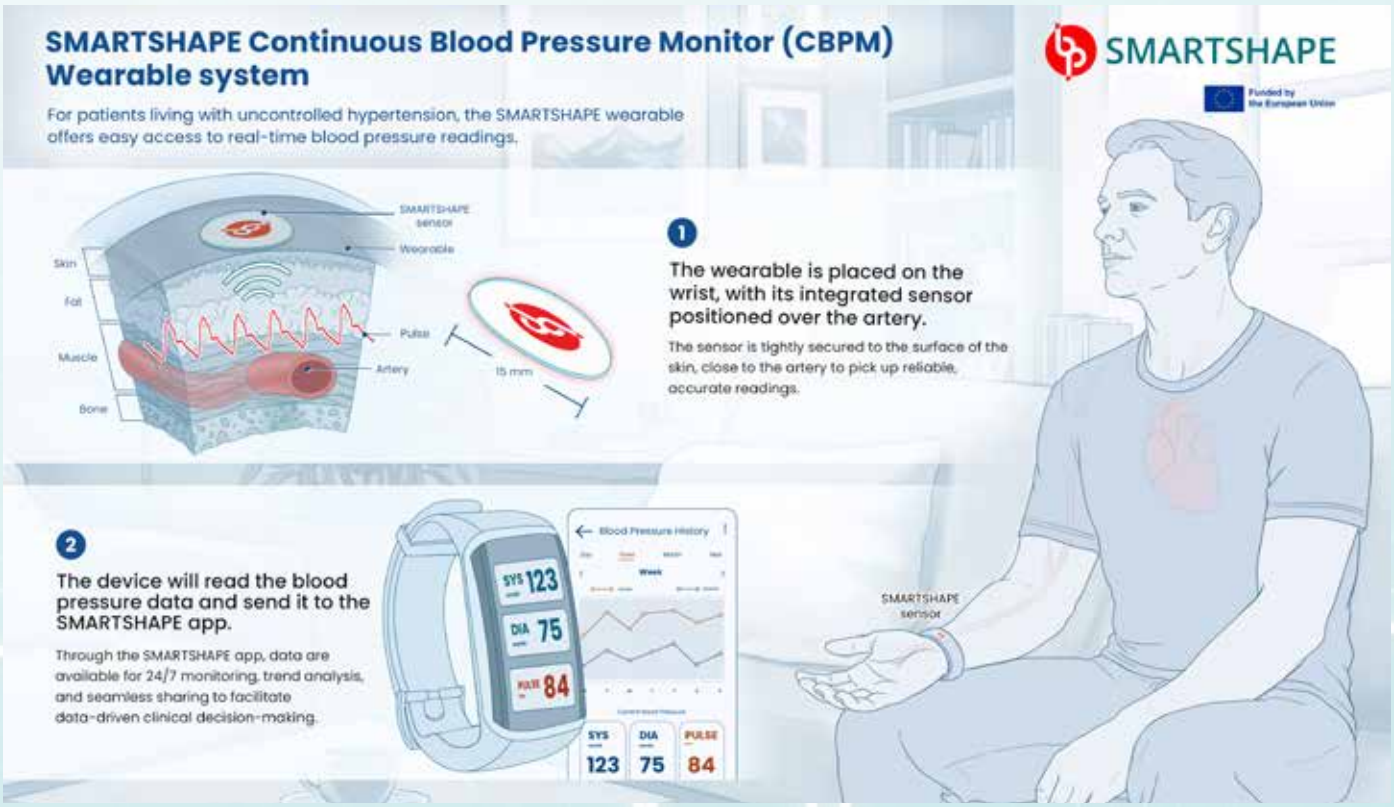
Galway has built the reader hardware and the software used to calibrate this information, filter, visualise, and capture data in the preclinical studies. Both now need to move towards product form: the reader must be miniaturised, while the software must itself satisfy medical-device requirements.

This chain from deformable material to clinical display is where SMARTSHAPE's interdisciplinary method matters most. A change made to ease fabrication can affect sensitivity; packaging can alter wireless performance; miniaturising the battery-powered external reader creates trade-offs between battery size, operating life and ergonomics; algorithms must distinguish physiology from movement; and the interface must transform an unprecedented volume of information into decisions rather than noise.

Patient and clinician co-design has therefore run alongside the engineering. Consortium partner Merakoi, which leads SMARTSHAPE's patient-engagement work, put together a patient panel spanning two distinct experiences: people who have suffered a stroke or heart attack, sometimes more than once, and people living with uncontrolled hypertension who have not experienced such an event. The first group knows the consequences first-hand, while the second is managing an invisible risk. Both have helped test the patient app and clinician dashboard while the system can still be shaped by their feedback.

Evidence beyond proof of concept

The strongest outcome so far is the performance of the sensor in an 84-day chronic preclinical study. Under anaesthesia, the team could compare SMARTSHAPE directly with an arterial



line. Connolly says the relevant continuous-monitoring criteria rightly demand stringent accuracy criteria to be met and quotes the relevant ISO80160-3 for continuous monitoring in which it demands that a new technology must achieve a mean difference of no more than 6mmHg and a first standard deviation of no more 10mmHg, when compared to the gold standard reference. On day 84, SMARTSHAPE achieved 2mmHg and 7mmHg respectively.

Those figures matter because they move the discussion from novelty to clinically relevant accuracy. "It ticks the box with respect to accuracy," Connolly says, while offering a simpler interpretation for non-specialists: "At that stage of preclinical testing, the sensor performed with accuracy comparable to the arterial-line reference used in surgery."

Parallel in-vitro work extended to 24 weeks, leaving the team with what he describes as "at least six months of accuracy data", rather than evidence of multi-year performance in patients.

The preclinical programme also produced usable data from an awake animal. "We even got nice daytime readings when the animal wasn't under anaesthesia," Connolly recalls. "You're kneeling down beside a sheep, holding the reader up to the neck and still getting nice pulse-wave data." It was, he says, "a nice bonus" and an indication that the system could operate beyond tightly controlled anaesthetised conditions.

"Following these results the question is no longer simply whether a passive implanted pressure sensor can produce an accurate signal," concludes Connolly. "It is whether the entire system can preserve that accuracy through human use, automated manufacture and the clinical evidence required for adoption."

To a care pathway

The patient research suggests that more data will only be valuable if it produces a more intelligible view of health. Participants responded particularly strongly to "time on target": the proportion of time an individual's blood pressure remains within an appropriate range, alongside the timing of excursions above or below it. This gives patients a trend that can be related to medication, sleep, exercise, work stress and daily events.

Kevin Michels-Kim, CEO and co-founder of Merakoi, sees particular value in the feedback loop that continuous monitoring could create for people taking two or three antihypertensive medicines. "Recording medication use becomes more meaningful when patients can relate it to changes in their readings, while context can be logged when something happens rather than reconstructed later in a paper diary," he says. "They've never had that view of their blood pressure, and now it's revealed through the app."

The other half of this feedback loop sits with clinicians. Neurologists, cardiologists, nurse practitioners and family-medicine specialists told the project team that they had "never had this kind of data available to them", Michels-Kim reports. "Home readings may be incomplete, while information from consumer devices is often discounted when its accuracy is uncertain. Given clinically reliable data, clinicians wanted both a population view, allowing hundreds or thousands of patients to be screened for alerts, and the ability to examine an individual's targets, comorbidities and treatment history."

Delivering the views clinicians want is not simply a matter of designing a better dashboard. The technology must also fit existing care pathways, digital infrastructure and reimbursement models. SMARTSHAPE Commercial Lead Ciarán Duffy sees adoption as a connected challenge encompassing clinical evidence, data governance and security, interoperability and workflow design. "We need to demonstrate a clear workflow in terms of patient integration, where it sits in the pathway," he says. "Implantation must be straightforward; the device must have a defined role in post-stroke, cardiovascular and step-down care; regulators and payers must see evidence of clinical value, and the dashboard must connect effectively with electronic health records and hospital systems."

Requirements for this will differ between the EU, UK and the US. Some clinicians told the team that they would initially tolerate a standalone system — and even copy data between platforms — if the information clearly improved care. Others saw integration as essential. Duffy argues that SMARTSHAPE cannot prove a dashboard in isolation and ask how to integrate it afterwards. "We have to have an interoperability strategy that is clear and that will work in the long term," he says.

Manufacturing presents the same need for whole-system thinking. The current sensors are assembled by hand with good consistency, but Connolly says the design has reached "the edge of how miniature we can go using human dexterity and further reduction will require high-tolerance automation.

Duffy's emphasis is design for manufacture: "A sensor that is quick, economical and robust to produce, with cost of goods compatible with reimbursement, yet still meeting its accuracy and reliability requirements."

The external reader adds its own scale-up questions around miniaturised electronics, wireless protocols, battery life, ergonomics, secure transmission and movement artefacts. Commercialisation is not therefore a simple transfer from a laboratory to production line. It is a redesign of sensor, reader, software and service around validated processes and viable unit economics.

Learning what not to assume

These implementation challenges have reinforced a broader lesson for SMARTSHAPE: technical performance alone will not determine whether the device is accepted in clinical practice. Patient co-design showed that willingness to use continuous

monitoring did not imply unrestricted consent to share the resulting data. Patients strongly supported sharing it with the clinician responsible for their care. "If it's sharing with my clinician, I'm 100 per cent on board," Michels-Kim reports. "They also made it clear that they wanted the device to be prescribed because, as one participant explained, 'then I know that it has a medical value'."

Sharing with relatives was different, with some participants not wanting a spouse or family member to have the same access, particularly when hypertension had not yet produced a visible medical event. "It's a very personal thing," Michels-Kim adds. "You need to give people autonomy over their data, and we have to build that into the solution."

The amount of information generated raises a further question. One neurologist asked how the project would use "millions or billions of data points" and envisaged researchers and artificial intelligence identifying patterns hidden by population-level thresholds. Connolly sees that potential but says the project is proceeding "with extreme caution".

"Patients might consent to their data being used to train machine-learning systems," he continues, "but not for clinical decisions".

Clinical validation has brought another unforeseen constraint in that a continuous monitor requires another continuous measurement as its reference, limiting potential study participants to people who already need an arterial line. "In order to validate our continuous technology, we need another reference continuous system in place already which can only be the invasive arterial line in the clinical setting," Connolly says. "The next trial must therefore resolve a practical and clinical-trial design challenge."

The next threshold

As the Horizon Europe project approaches its conclusion, SMARTSHAPE is seeking the resources needed to continue development. Duffy says the team "has already secured additional funding to develop the technology". The next phase will include manufacturing at scale, further miniaturisation, regulatory work and clinical trials.

Within five years, Connolly wants to see the technology operating in patients, generating real-world evidence and overcoming its principal manufacturing challenges. "The most important thing for the project," he says, "is seeing it working in patients and having a positive impact."

Duffy hopes that clinical adoption will eventually reduce the healthcare costs associated with hypertension, although he stresses that "the whole clinical utility question is paramount". Michels-Kim envisages hypertension being managed through "evidence and data rather than the one-off readings" used today, with fewer patients falling through gaps between appointments.

Continuous monitoring could reveal nocturnal hypertension or significant changes that intermittent measurements miss, but its preventive value still has to be demonstrated. "We can't claim it yet because we'd have to run studies on that," Michels-Kim cautions. That is SMARTSHAPE's next test: showing that accurate, beat-by-beat information can improve clinical decisions and patient outcomes.

PROJECT INFORMATION

Project Title
SMARTSHAPE - Development of a smart shape-memory sensor for biological pressure monitoring applications

Project Objective
By developing a medical device that allows patients to monitor their blood pressure levels 24/7, SMARTSHAPE aims to help improve healthcare for people living with hypertension and support the development of personalised – and more effective – diagnostic and treatment plans. SMARTSHAPE delivers continuous, clinically accurate blood pressure monitoring across two complementary form factors – a subcutaneous insertable system for high-risk patients in clinical or step-down care, and a wearable system for broader ambulatory use.

Project Duration and Timing
Start date: 1 December 2022. End date: 30 November 2026

Project Funding
€ 3 218 279,51. HORIZON-CL4-2022-RESILIENCE-01-13 - Smart and multifunctional biomaterials for health innovations (RIA)

Project Partners
University of Galway (coordinator), TSQ Global S.R.L., Scinnovation s.r.o., Explora SRL, Merakoi AG, University of Birmingham, Catalyze Group

MAIN CONTACT



Name
Ciaran Duffy
Commercial Lead
Ciaran.duffy@universityofgalway.ie



David Connolly
Cardiovascular Research Programme Manager
David.connolly@universityofgalway.ie

Web address
www.smart-shape.eu

SMARTSHAPE

This project has received funding from the European Union's Horizon Europe research and innovation program under the grant agreement **No 101092242**.

INSIGHT FEATURE

Helping manufacturing research reach the factory floor

Europe's manufacturing research is developing technologies that could improve productivity, reduce waste and strengthen industrial competitiveness. Realising that potential requires effective communication, informed industrial engagement and sustained support for adoption. Drawing on its close relationship with EFFRA and years of experience within manufacturing projects, Insight Media's **Insight Access** service gives projects, research organisations and industry flexible access to the expertise needed to connect research with practical use. **Sam Davis**, Commercial Director at Insight Media explains more.

A successful industrial demonstration gives manufacturers valuable evidence about a technology's potential. Turning that evidence into wider adoption requires further conversations: how the solution fits existing equipment, what validation remains necessary, who will support implementation and where the business opportunity lies.

Those conversations should develop alongside the research. Understanding industrial needs can help researchers refine their approach, identify applications and present results in a form that potential users can assess. Maintaining that exchange takes time, specialist skills and a clear understanding of the audience.

For an individual project, the challenge may be delivering an ambitious communication and dissemination programme. For a research organisation managing many projects, it may be coordinating competing demands across a small internal team. For an industrial partner, it may be explaining a new technology to prospective customers or finding collaborators for its next development stage. Insight Access is designed to support all these needs.

Understanding the manufacturing community

At Insight Media, we have worked as a communication, dissemination and exploitation partner within manufacturing projects over many years. That experience has helped us understand both the research process and the practical questions that influence industrial uptake.

Our close relationship with the European Factories of the Future Research Association, EFFRA, strengthens that understanding. Through our work with the manufacturing community and publications including EFFRA's anniversary edition, we explore how research contributes to industrial priorities and where collaboration can help technologies progress.

The projects featured in this magazine reflect the breadth of that activity, from advanced materials and production processes to robotics, sensing and distributed intelligence. Each has different potential uses and audiences to reach, so it is important to tailor communications to explain what each project offers and why it matters to those who could use its specific results.

Flexible expertise

Insight Access gives organisations access to Insight Media's communication, dissemination, stakeholder engagement and exploitation expertise as an extension of their teams. Support can develop with their needs, from early planning and audience development through demonstrations, results dissemination and commercial engagement.

Projects can draw on additional capacity to deliver their work plans, explain technical progress and engage stakeholders throughout implementation. Existing strategies can be refined as results emerge and opportunities become clearer. That can be particularly important at the end of a project, when audience engagement is so vital.

For universities, research institutes and technology organisations with multiple projects, Insight Access provides continuity across a changing portfolio. It can support busy periods, bring specialist skills to particular activities and help teams connect related work with relevant audiences. Knowledge developed through one assignment can inform subsequent activity, reducing the need to brief a new supplier each time.

Industry can use the same expertise to explain research-led products and processes, communicate evidence from trials and engage prospective customers or development partners. Working with technical teams, we help translate complex capabilities into information that supports informed business decisions.

Making results useful

Our films, highlighted as exemplary by reviewers and European Commission observers, alongside brochures, magazines, policy briefs, websites and technical content give this work practical substance.

The choice of format depends on what a project is trying to achieve. A film can show how a process operates under industrial conditions. A technical brochure can explain performance, integration requirements and potential applications. A policy brief can set out the implications of research for decision-makers.

Working closely with researchers and industrial partners, we establish what has been demonstrated, what remains under development and what the audience needs to understand. That determines the type of engagement activity needed, and that clarity helps potential users judge relevance and identify a sensible next step.

Developing your audience

Effective content needs to reach people with a reason to act on it. A production manager assessing an inspection system, for example, needs different information from a machinery supplier considering integration or a research organisation seeking a development partnership. We segment audiences by sector, application, role and interest, then build engagement around those needs.



Sharing relevant evidence and inviting feedback helps identify genuine interest. Following up can develop warm leads: organisations with an expressed need, a suitable application or an interest in trials, licensing, integration or further collaboration. This audience development has value throughout a project or organisation's work. It can inform research priorities, support demonstrations and strengthen the relationships on which future business depends.

Strengthening exploitation

Of course, projects already have exploitation plans set out, developed in their proposals and part of their ongoing work plans. Insight Access helps partners strengthen and implement those commitments around the results achieved, the interest generated and the development still required.

This becomes particularly valuable near project completion, when consortium resources and responsibilities are changing. A promising industrial conversation may need further evidence, a meeting with the result owner or continued contact while a prospective user assesses its options.

Clear post-project planning establishes who will take those opportunities forward and how engagement will be maintained. Insight Access can provide ongoing communication and dissemination support, keeping useful information accessible and helping industrial partners pursue relevant contacts beyond the funding period.

That work supports the technical and commercial decisions of the organisations responsible for the results. It gives them a better-informed audience and a stronger basis for discussions about adoption.

Partnership for manufacturing success

Europe's manufacturing ambitions depend on the ability to connect research expertise with industrial capability. Communication, dissemination and exploitation support help those connections develop throughout the journey from an initial idea to practical application.

Insight Access makes that expertise available in a flexible form: supporting individual projects, adding capacity across research portfolios and helping industry engage with the people who can use its innovations.

Our aim is to build lasting working relationships around manufacturing research, giving organisations the support to explain their achievements, understand their audiences and take the next step. Talk to us and see what we can do for your results.

sam@insightm.co.uk

CONNECTED INTELLIGENCE

EXTRA-BRAIN: Thinking differently for a new kind of intelligence

Artificial intelligence can achieve remarkable results when data is plentiful, computing power abundant and conditions remain stable. The real world is rarely so accommodating. **Pawel Herman,** professor at KTH Royal Institute of Technology in Stockholm and coordinator of EXTRA-BRAIN, explains how the EU-funded project is drawing on the brain's architecture and learning mechanisms to explore AI that can adapt to change, use resources more efficiently and ultimately achieve more beyond the laboratory.

The most exciting promise of artificial intelligence is not simply that machines will become better at repeating tasks we already know how to automate. It is that they may help people and organisations act intelligently when information is incomplete, conditions are changing, and the right answer cannot be anticipated in advance. In a factory, that could mean a robot that adapts to a new production sequence or monitoring systems that identify emerging defects before they affect production. At the edge of a communications network, it might mean allocating resources without continually consulting a distant data centre. In an emergency, it could mean finding a person in a place no person or machine has ever seen before.

Yet much of today's AI delivers these impressive performances through sheer scale, with huge models, more training data and formidable computing power. That formula does work exceptionally well, but it also creates a demanding industrial proposition. Data must be collected, labelled, governed, stored and moved; models may take days to retrain; and the energy and hardware they require can put them beyond the capabilities of small devices.

EXTRA-BRAIN is a Horizon Europe project that starts with a different proposition. Rather than considering how far conventional deep learning can be scaled, it investigates what AI can gain by returning to the source that first inspired neural networks: biological intelligence, the human brain.

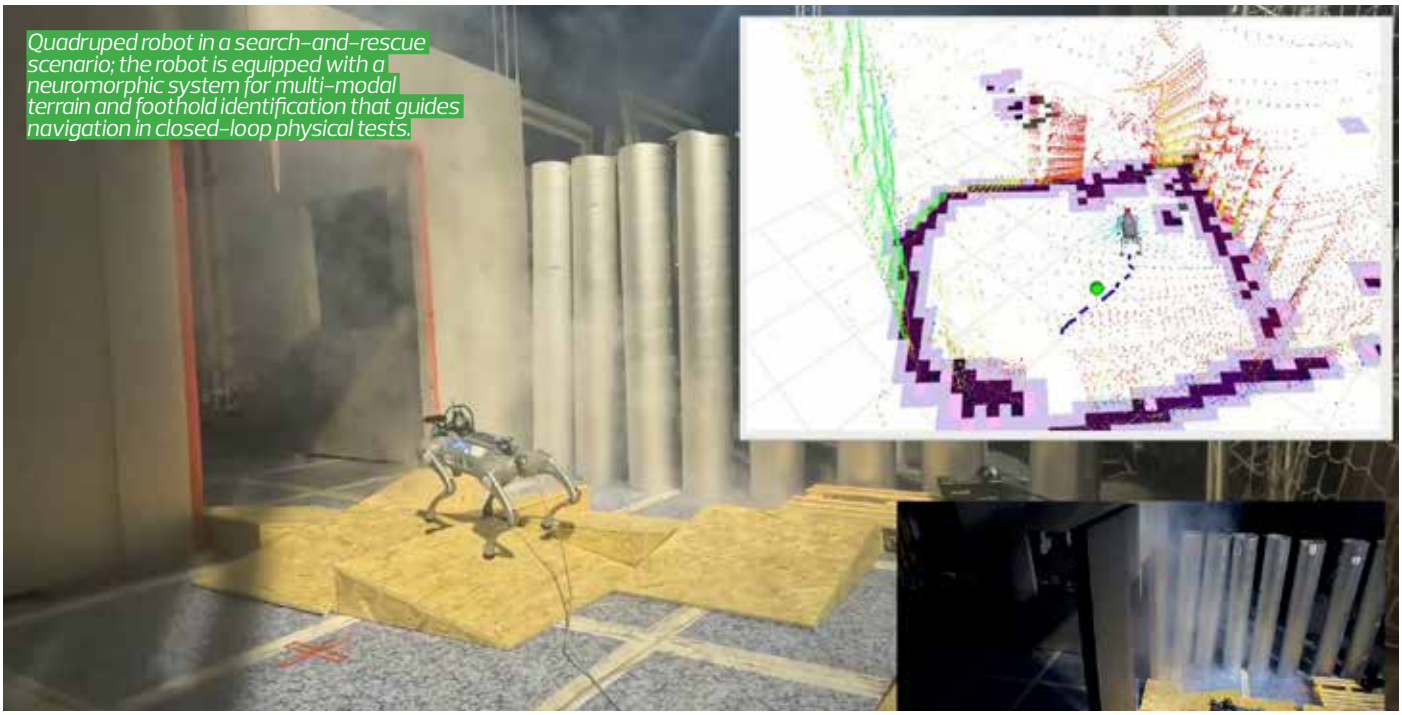
Its models draw on the organisation and learning mechanisms of the brain and are developed alongside neuromorphic hardware, in which computation is engineered around similarly distributed, event-driven principles.

"The distinct feature of this project is that our approach builds on what we would call brain-like computing technology," explains Pawel Herman. "The mainstream builds very much on deep learning, whereas we rely on neural networks inspired by what we know about neural information processing in the brain. That is coupled to neuromorphic computing through the hardware."

Herman is not proposing that every conventional model should be replaced and he is notably pragmatic about where brain-inspired methods may offer their greatest value. "If you have a huge amount of well-prepared, well-labelled data and huge energy resources, then I would recommend that you go with standard mainstream AI," he says. "The moment you deviate from that scenario, this is where I would start looking for a neuromorphic solution."

Of course, deviation is common on the factory floor and throughout the real world of industry. Sensors become noisy, communications fail, new products enter a production line and operating conditions drift. Industrial AI must therefore be able to work beyond the static conditions for which it was trained, just as people routinely do.

This is the practical problem EXTRA-BRAIN is designed to address with its brain-like networks, which are intended to cope with incomplete or unlabelled data, adapt through local learning rather than wholesale retraining, and indicate when an output is uncertain. "This is where the concept of trustworthiness begins for us," Herman explains. "It is about AI remaining useful when conditions change, rather than performing well only when the data and environment resemble those encountered during training. Crucially, we aim to achieve that adaptability



Quadruped robot in a search-and-rescue scenario; the robot is equipped with a neuromorphic system for multi-modal terrain and foothold identification that guides navigation in closed-loop physical tests.

with less data and lower computing demands. However, when eventually reliable predictions cannot be made in dynamically changing scenarios any longer, trustworthy AI should clearly communicate its lack of confidence and adapt by acquiring new knowledge."

The significance of that efficiency extends beyond technical performance as it could determine how widely advanced AI can be used. A system that learns from smaller quantities of data and operates with less computing power requires less infrastructure to develop, adapt and deploy. "That makes advanced capability more practical beyond organisations with the largest datasets and computing budgets, while supporting its deployment on robots, sensors and other edge devices where power, memory and connectivity are constrained," adds Herman. "Resource efficiency is therefore not simply an engineering advantage; it could extend capable AI into more of the places where industry actually needs it."

There is also an economic and democratic dimension; useful AI depends so much on the largest proprietary datasets and the most expensive computing infrastructure, meaning advanced capability will remain concentrated among a small number of organisations. Systems able to learn from fewer examples, including unlabelled observations and perceptions gathered during operation, could bring sophisticated intelligence within reach of smaller manufacturers and public services. They could also keep sensitive information close to where it is created, reducing the need to transmit every signal to a central platform, potentially reducing the exposure of sensitive data.

What the brain can teach a computer

The project is not attempting to recreate the human brain, to reproduce a biological copy neuron by neuron. Its aim is purely functional: to identify architectural and learning principles that can make an engineered system behave more usefully. One starting point is connectivity. Conventional artificial neural networks are often dense, with many computational units linked across layers. The brain combines feed-forward, recurrent and backward connections in recognisable patterns, while remaining remarkably sparse.

"There is a structure, and that structure makes sense," says Herman. "We try to learn a lot from neuroanatomy, specifically in the cortex, where much of the cognitive computation takes place. What comes with that is very sparse connectivity, and this is where we save energy, but also where we gain robustness.

"Fewer active connections mean fewer operations and less costly communication, but sparsity also reduces the risk that one damaged component will bring down the whole system."

The second lesson concerns how an AI system learns. Mainstream deep learning typically adjusts the model by measuring a global error and propagating corrections throughout the network. EXTRA-BRAIN investigates local learning, including Bayesian-Hebbian mechanisms inspired by the way neighbouring biological neurons strengthen or weaken their connections. Because learning is distributed, the system can update the parts that need to change instead of recalculating everything. This creates a route towards continuous or lifelong learning at a much lower computing cost.

"In classical mainstream networks, in order to make any changes, you need to change all the weights. It is a heavy process," Herman says. "In our case, you make adjustments because learning happens locally; it is distributed."

"The brain also does not learn indiscriminately and global modulation tells it when an event is important enough to alter memory and when information can simply be processed," he continues. "EXTRA-BRAIN is building a comparable distinction between episodes of learning and episodes of computation. In practical terms, the network does not have to alter itself whenever it receives new input: it can process routine information using what it already knows and make targeted adjustments only when something sufficiently novel or important appears."

Memory itself is treated differently too. In conventional computing, information repeatedly travels between processing and storage. In the brain, memory and computation are intertwined. "For us, memory, like in the brain, is the place where you actually perform computations," Herman explains. "Co-locating the two makes memory an active part of the processing, rather than simply a store for what the system has learned."

Another biological reference point for the project is the way a neuron in the brain communicates through discrete events, known as spiking. This is an all-or-nothing form of communication. A neuron that remains quiet sends nothing to its neighbours; only when it becomes active does it transmit a small unit of information. That selectivity is attractive because, in principle, the network communicates, and uses energy for that communication, only when there is information to pass on. For AI designed to operate on compact devices with limited power, the potential benefit could be significant.

Herman cautions, however, against assuming that spiking is inherently energy efficient and fast. Reliable information processing rarely depends on a single spike; it normally requires a series of them, and generating those additional events consumes energy and takes time. EXTRA-BRAIN's investigations have therefore placed the speed and energy efficiency of spiking and non-spiking versions of its networks at a similar level.

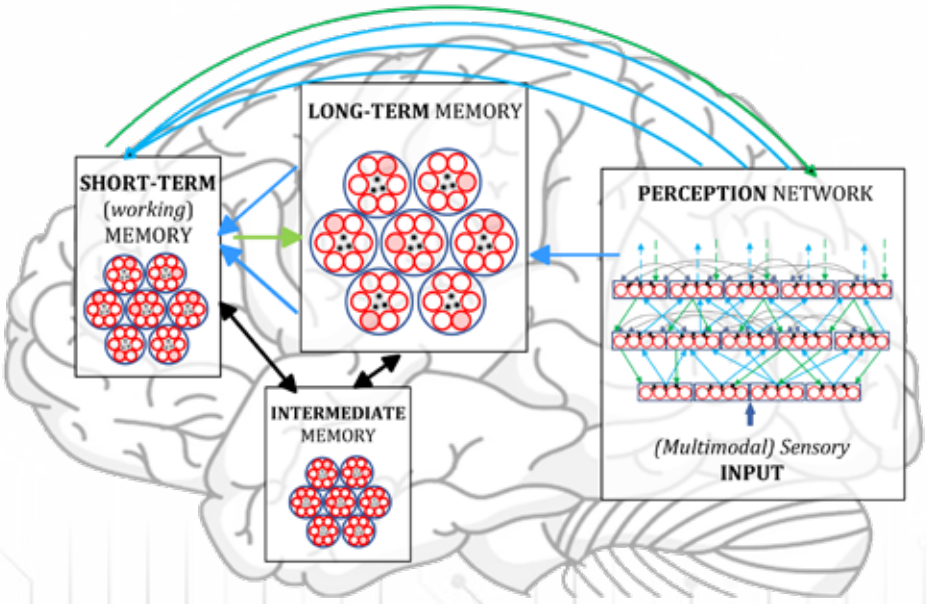
The greater opportunity may lie in matching spiking algorithms with hardware designed to handle this event-based communication efficiently. The value comes not from the spike alone, but from the relationship between the network, the chip on which it runs and the application it must serve. "Only with that co-design can you squeeze out from the algorithms what is most beneficial for applications," Herman says. "The algorithm, hardware and intended use must therefore be shaped together rather than optimised in isolation."

Trust when conditions change

For industrial users, reliability is usually discussed in terms of accurate and repeatable performance. EXTRA-BRAIN adds another dimension to this with "the ability to remain useful with poor, missing or uncertain information". A brain-like network can work with unlabelled data, fuse different kinds of input and tolerate noisy sensors or even the loss of computational nodes.

"Because its knowledge and processing are spread across many interconnected units, the system does not depend on every component working perfectly," Herman explains further. "If a sensor, connection or computational node is lost, the remaining network can continue to produce useful results, although its performance may decline. Rather than failing outright, it degrades gradually and can potentially adapt around the problem.

"A distinctive feature of reliability would be resistance to poor data, lacking data and uncertainty," he continues. "If you switch off some nodes, you can quite quickly readapt or reconfigure the system without spending a lot of time or requiring a lot of data."



Brain-like neural network architecture that builds on self-organised feedforward pathways for unsupervised learning of input representations and recurrent networks for gestalt perception and associative memory-related pattern processing.

The probabilistic representation of information also allows the model to report how certain it is about an output. That matters whenever a human decision-maker must distinguish a confident recommendation from an educated guess.

Technical resilience alone, however, does not earn human trust and EXTRA-BRAIN, therefore, complements technical assessment with a cycle of stakeholder-centred validation workshops, presenting its emerging technologies to the people who might use, develop or be affected by them and redesigning them in response to what those stakeholders find useful, unclear or missing.

The priorities differ according to the application, with frontline users perhaps caring more about whether a system provides the functions they need under difficult conditions, while technical specialists may require a closer understanding of the algorithms and their deployment constraints. An ethical adviser helps the consortium interpret both legal requirements and the less easily codified expectations surrounding human-centred AI. This activity also reveals where dependable performance is not enough and users need a clearer understanding of how the system reaches its outputs.

That need for users to understand how a system reaches its outputs brings explainability into focus and Herman regards this as one of EXTRA-BRAIN's key challenges. "Because knowledge is distributed across many interacting units, a brain-like network does not naturally provide a simple account of how a particular result was formed," he explains. "At the moment, we treat this as a bit of a grey box that offers some, still limited, insights into salient data features. The project's explainability work is therefore examining how the networks can be more systematically used to make their outputs more interpretable."

"For the time being, I see a certain weakness, though there are emerging ideas about how to turn these into fully-fledged explainable systems," he adds.

"Trustworthy AI is not a label applied after development; it requires continuing attention to performance, evidence, human oversight and the limits of what a system can explain."

Explainability is only one part of a wider engineering balancing act the project is considering. Every deployment involves choices about the accuracy and adaptability a system can deliver within the available data, computing power and energy. Those constraints become particularly visible at the edge: a network designed to operate on a sensor using a watt or less must remain small, which inevitably limits its performance. "The trade-offs are inevitable: size, scale, energy and performance," Herman acknowledges.

EXTRA-BRAIN's proposition is not that brain-like models remove those compromises but that their local learning, adaptability and ability to work with imperfect data may give designers more freedom in how they balance them, particularly when robustness under real operating conditions matters as much as peak performance on ideal data. The project's three demonstrators provide the practical settings in which that balancing act can be tested.

Three proving grounds

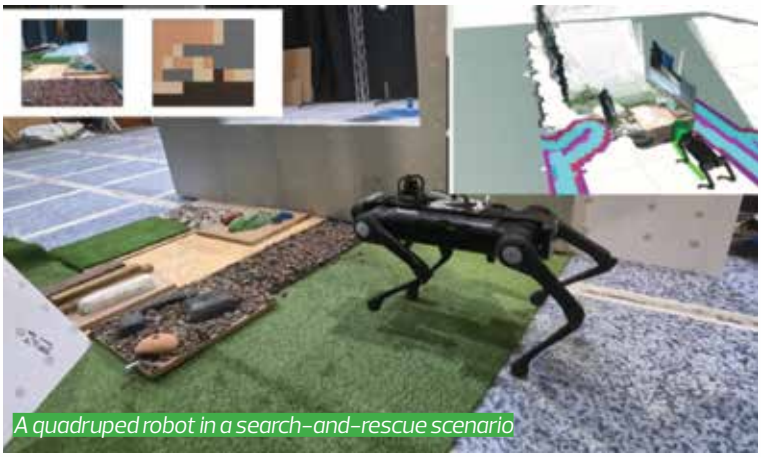
The project has three deliberately contrasting demonstrators: search-and-rescue robotics, data-efficient digital finance and the allocation of telecommunications edge resources. Together they ask whether the same brain-inspired computational principles and brain-like network motifs can support perception and action in the physical world, interpret uncertain time-dependent information and solve a complex optimisation problem.

The robotics case offers the clearest picture of AI operating beyond a controlled environment. Existing robots are being equipped with neuromorphic perception and learning so they can search for victims trapped in disaster areas. Recognition is only the first step. The intended system must observe its surroundings, keep relevant information in memory, decide where to look next and learn while it moves. "They should not only recognise objects and find potential victims; they perform a more active search," Herman says. "We will equip them with memory function so that they could navigate in the environment, observe things, memorise and prioritise what is important and then act on this."

Running intelligence on the robot itself is an important demonstration of the project's potential. "A machine sent into an earthquake or fire cannot assume a stable cable, cloud link or nearby computer," says Herman. "We are working to place the neuromorphic hardware physically on the robot so that perception and decision-making can happen in the field with a lower energy footprint."

Early perception tests compare well with classical approaches while using less data, but the more revealing integration of memory, adaptation and decision-making is the next step. The objective, Herman stresses, is "not so much to demonstrate that it works better, but that it offers more functionality than the classical approaches".

Digital finance tests a different strength as its streams of market data are time-dependent, heterogeneous and inevitably uncertain. EXTRA-BRAIN's networks can fuse those streams, learn from relatively small datasets and produce a probability estimate as part of a single run. That last capability proved especially important to financial advisers. "What is really valuable is not only the prediction accuracy, but the uncertainty around this," Herman says. "Conventional methods may need to run repeatedly to simulate randomness; a probabilistic brain-like network can handle and describe uncertainty directly, reducing computation."



A quadruped robot in a search-and-rescue scenario

The telecommunications demonstrator addresses the optimal allocation of edge-computing resources, a problem in which many possible arrangements must be evaluated against multiple constraints, rather than a single image or signal being classified. Memory-based computation makes the project's networks a promising match. "The work is not complete, and effective learning within the optimisation process remains challenging," says Herman. "But the use case broadens the evidence beyond tasks for which neural networks are already a familiar tool."

Across all three cases, the practical themes are closely aligned with manufacturing and the requirements AI must meet on the factory floor. Production systems bring together diverse sensors, constrained edge devices, intermittent connections and processes that change over time. Models that can learn locally, operate with limited data and recover when a component fails could support robots that are easier to repurpose, inspection systems that remain useful as image quality changes, and distributed control that continues operating when a cloud service is unavailable.

The lesson is not that one artificial brain will run an entire factory, but that specialised, cooperative intelligences may divide responsibilities and work together. Herman instead envisages "functional systems of brains that can collaborate, cooperate, share responsibilities and do things together."

That additional functionality is only the beginning of the ambition, and for Herman that means EXTRA-BRAIN making neuromorphic, brain-like computing more visible to industry. "We want to demonstrate that it could become a competitive European capability," he says. "If you can open up the industrial world's eyes to this neuromorphic, brain-like computation, which could offer Europe a competitive advantage, perhaps there will be more funding and a move towards the mainstream," he continues. "That's something which I see as a mission for EXTRA-BRAIN."

Behind that mission lies a still larger question. Artificial neural networks began with inspiration from the brain but progressively moved away from its architecture and learning mechanisms. EXTRA-BRAIN is now exploring what might become possible by returning to those biological foundations – not to reproduce the brain mechanically, but to draw more consistently on the principles that make it so adaptable, selective and efficient. For Herman, the fascination lies in "closing that gap, really seeing what we can actually achieve if we are more consistent in sticking to how the brain works."

PROJECT INFORMATION

Project Title
Explainable Trustworthy brain-like AI for Data Intensive Applications (EXTRA-BRAIN)

Project Objective
The EU-funded EXTRA-BRAIN project aims to develop flexible, neuromorphic hardware friendly, energy efficient, and trustworthy AI solutions based on brain-like neural networks.

Project Duration and Timing
01.01.2024 - 28.02.2027

Project Funding
4 994 313 EUR; HORIZON Europe

Project Partners
KTH Royal Institute of Technology (Coordinator) + 11 project partners:

Universidad Politecnica de Madrid

Centre for Research and Technologies

Ericsson

GFT Italy

Unparallel Innovation

Brimatech Services

Xgility

i4RI

ICE - Information Catalyst

Nerion

ExpertAI - LUX

MAIN CONTACT

Name
Pawel Herman



Phone
+46 8 790 65 13

Email
paherman@kth.se

Web address
www.extra-brain.eu

This project has received funding from the European Union's Horizon Europe research and innovation program under the grant agreement **No 101135809**.

FREE SUBSCRIPTION

Apply for your free subscription to **Projects Magazine**

INSIGHT MEDIA

Projects

Showcasing research and innovation in Europe

Projects is Europe's leading publication focussing on research and innovation. The magazine has established itself at the forefront of innovation by providing a critical platform for the dissemination of projects' results as well as being the natural home for the debate and analysis of the latest issues relating to policy and funding. Projects also highlights the success stories that drive Europe's innovation and is a must-read for all those involved in research, from academia, industry and the public sector.



Insight Media Group Ltd.
King Square, Bristol, BS2 8JH, UK
Prins Hendrikkade 21-E, 1012TL, Amsterdam, Netherlands
Tel: +44 (0)117 248 2755

CONNECTED INTELLIGENCE

OpenSwarm: Intelligence in numbers

Thousands of small, low-power devices already monitor the physical world. OpenSwarm set out to make them interpret it, collaborate and act. **Micael Couceiro**, CEO and co-founder of Ingeniarius and **Thomas Watteyne**, Senior Research Director at Inria and the project's scientific coordinator, explain how the €5.32 million Horizon Europe project combined dependable networking, constrained AI and energy-aware swarm programming – then put the resulting architecture to work across large-scale testbeds and five proofs-of-concept.



The OpenSwarm testbed of 725 DotBots

For years, the dominant model for low-power wireless systems has been simple: a sensor takes a measurement, sends it to a server and waits for a central application to interpret it and then what to do with that information.

OpenSwarm – Orchestration and Programming ENergy-aware and collaborative Swarms With AI-powered Reliable Methods – is a project that started with a different proposition: what if small, energy-constrained nodes could perceive their surroundings, interpret selected information locally and coordinate with neighbouring devices? Useful behaviour would no longer come only from a central server. It would emerge from the interaction of many participants, each making a limited but meaningful contribution to a shared mission.

"The fundamental difference is where intelligence and responsibility reside," says Micael Couceiro, CEO of Ingeniarius, a Portuguese robotics company and OpenSwarm partner. "In a conventional monitoring network, devices are often essentially data sources: they acquire measurements, transmit them to a central server, and wait for something else to interpret the information and decide what to do. A collaborative swarm is different. Each node has some degree of perception, local intelligence and autonomy, and the useful behaviour emerges from the interaction between nodes. Where an IoT network connects devices, a collaborative swarm turns those devices into cooperating agents."

The opportunity is clearest wherever the physical world is too large, dynamic or inaccessible for a single device or centralised system to handle efficiently. "We see this

very clearly in field robotics," Couceiro says. "Forests are a good example: connectivity is imperfect, Global Navigation Satellite System (GNSS) can degrade below the canopy, the terrain is unstructured and there may be no reliable cloud connection. A robot must therefore retain enough intelligence locally to continue operating even when communications deteriorate."

Marine monitoring and agriculture share the same characteristics – and so, at a different scale, does a factory floor containing thousands of machines and mobile robots. Sending every raw measurement to one central system is unnecessary, expensive and potentially fragile. Local intelligence can shorten the path from perception to action, while the network reserves its capacity for decisions that genuinely need a system-wide view.

OpenSwarm's ambition was to show that intelligence could be distributed across such environments without sacrificing



A Renewable Energy Community proof-of-concept demonstration at the Swarm Workshop in Brussels, 26–27 November 2025

reliability, using simulations, large-scale testbeds and field trials on real hardware. Distributed intelligence promises faster decisions, greater resilience and lower network traffic: instead of continually transmitting raw point clouds or multi-channel acoustic recordings, nodes can exchange only the features, events or decisions the wider swarm needs. It also turns energy from a battery specification into a system-wide design problem.

When scale changes the problem

Turning that ambition into a dependable working system presented challenges at every level, beginning with energy. OpenSwarm had to treat it as a system-wide design constraint rather than simply a battery specification. "A traditional approach might consider battery level only as a safety threshold: below a certain voltage, the robot returns home," Couceiro says. "In a collaborative autonomous system, that is too simplistic. Energy influences task allocation, route selection, mission duration, communication strategy and even whether a robot should continue a task or transfer responsibility to another member of the swarm."

The project investigated this directly in its forestry-drone work through an energy-aware capability developed with imec. The research examined how predicted consumption, based on trajectory, velocity, altitude and mission conditions, could become an input to a drone's autonomy rather than a threshold checked at the end of the decision chain. At swarm level, Couceiro explains: "The objective is not to maximise the endurance of every individual node independently. It is to use the available energy across the collective in a way that maximises mission effectiveness while maintaining safety."

A second practical challenge was heterogeneity. "One of OpenSwarm's strengths is recognising that a swarm architecture cannot be used in an industrial setting if it assumes homogeneous research hardware and unrestricted physical access to every node," says Thomas Watteyne, the project's scientific coordinator and Senior Research Director at Inria. The architecture therefore had to function across a huge range of platforms, from severely constrained microcontrollers in KiloBots and DotBots to ROS-enabled drones with considerably more processing power.

Watteyne notes that variation existed even within a single platform. With 2,000 motors running across the DotBot fleet, small differences in motor speed meant localisation and control algorithms could not simply be tuned once and copied across the swarm. The team instead developed self-adaptive methods that combined noisy sensor measurements and compensated for differences between individual devices.

Secure enrolment, firmware management and over-the-air updates presented a related challenge. Couceiro's later involvement in the marine-monitoring proof of concept made the practical importance of remote management particularly clear: retrieving a buoy whenever software or AI parameters needed to change would make a persistent deployment prohibitively expensive.

A DotBot




The OpenSwarm Consortium Meeting 2026 in Erlangen, Germany

Scale introduced problems that remain largely invisible in demonstrations involving only a few devices. "Many problems that appear negligible with five robots become dominant with hundreds," Watteyne explains. "How do you configure hundreds of nodes consistently? How do you deploy an update? How do you know which node failed? How do you collect logs? How do you reproduce an experiment? How do you initialise and recover the system without manually touching every device?"

Observability and automation therefore become part of the swarm technology rather than supporting administration. Couceiro encountered the physical equivalent in the forestry trials, where increasing the fleet from three UAVs to six and finally nine multiplied the demands associated with launch and recovery, radio-frequency coordination, battery management, field logistics and safety. The team supplemented the nine-aircraft deployment with simulations involving 30 UAVs, allowing algorithmic scalability to be explored without obscuring the operational constraints of working with real machines. As Couceiro puts it, "scaling a swarm is a systems-engineering problem, not simply a matter of increasing a parameter called 'number of robots'."

That is why, Watteyne explains, OpenSwarm combined large-scale testbeds with field proofs-of-concept (PoCs). The testbeds allowed technologies to be stressed systematically across hundreds of nodes; field deployments then introduced conditions that laboratories and simulations inevitably simplify, from weather and radiofrequency (RF) interference to physical damage, battery logistics, deployment time and imperfect hardware. "Either one alone would have been insufficient," he says.

Couceiro points to the marine trials as an example of a different but equally valuable kind of evidence. Distributed hydrophones on floating platforms had to remain precisely synchronised despite limited communications and short field-test windows. "The third proof of concept (PoC3) exposed limitations, including communication range and remaining embedded-deployment work," he says. "These types of results are equally valuable, as confidence in a framework comes not only from demonstrating what works, but also from understanding clearly where its maturity boundary lies."

Three pillars, two testbeds and five proving grounds

Running from January 2023 to April 2026, the €5.3 million, ten-partner project was coordinated by Inria and included Ingeniarius, Analog Devices, imec, KU Leuven, three Siemens organisations, the University of Sheffield and Wattson Elements. It organised its research around three interdependent pillars: efficient networking and management of smart nodes; collaborative, energy-aware artificial intelligence; and energy-aware swarm programming. The three pillars were designed to work together rather than in isolation.

The first pillar advanced low-power networking, with Time-Slotted Channel Hopping (TSCH) scheduling transmissions and switching radio channels, to make wireless links more deterministic and resilient. OpenSwarm also developed Mari, a custom link layer for mobile robot swarms. Watteyne says that "up to 150 robots can communicate with end-to-end reliability reaching 100%, and latency below the blink of an eye." For robotics integrators such as Ingeniarius, this means the network moves from networking designed mainly for static sensing towards networking capable of supporting collaborative and potentially mobile agents."

A parallel collaboration between Ingeniarius and Analog Devices integrated SmartMesh into the ROS (Robot Operating System) framework and tested how essential coordination data could keep flowing when the communication environment turned hostile. It provided proof, Watteyne says, "that reliable collaboration does


OpenSwarm's proof-of-concept project for monitoring ocean noise pollution using a low-power distributed network of hydrophone-equipped smart buoys

not always mean communicating more. Often it means identifying what information is essential and guaranteeing that this information can still flow under constrained conditions."

The second pillar addressed what should be computed where. "Not every node has to execute the complete AI pipeline, and not every piece of raw data has to leave the node," Watteyne explains. In the marine proof of concept, no single hydrophone could provide the spatial information needed for acoustic localisation, but continuously transmitting full raw audio from every buoy was prohibitively expensive in bandwidth and energy. In PoC3, OpenSwarm explored architectures in which lightweight processing is performed locally and only compact features or inference-relevant information are transferred for higher-level multi-node fusion, reducing bandwidth and energy requirements compared with continuously streaming raw audio. Watteyne stresses that this approach must be designed alongside the communications and energy architecture, "because the best algorithm in isolation is not necessarily the best algorithm for a field system."

The third pillar set out to change the level at which robotic software is written. "Robotic software is still written primarily from the perspective of the individual machine: go to this position, execute this trajectory, acquire this sensor data," Couceiro says. "Swarm programming requires us to move one level above: describe what the collective should achieve, under which constraints, while allowing the implementation to distribute that work among a changing group of agents." Making energy an explicit constraint matters in practice: if one of nine drones mapping a forest is burning through its battery faster because its assigned area is denser, an energy-aware swarm should be able to use that prediction to reconsider the mission. For example, by redistributing part of the task, adapting the trajectory or initiating safe homing before a critical threshold is reached.

To test the three pillars together, the consortium used two platforms: the KiloBot platform operated by the University of Sheffield and the DotBot platform developed by Inria, with the large-scale OpenSwarm testbed deployed at Analog Devices in Limerick, Ireland. The DotBots also supported an "Agentic testbed harness" that converts an operator's natural-language description of the desired swarm behaviour into code that can be compiled and loaded onto the robots over the air.

Outcomes

The clearest evidence that this approach works is the scale at which it has now been demonstrated. Mari demonstrated communication with up to 150 robots per gateway, reaching up to 100 per cent end-to-end reliability in the evaluated configurations; the consortium manufactured 1,028 DotBots and validated the OpenSwarm stack with up to 725 DotBots in the final large-scale experiments; and the KiloBot testbed at Sheffield provided a second large-scale platform built on entirely different hardware.

The forestry proof of concept, PoC2, provides particularly granular results. "All agreed final KPIs were achieved," Couceiro says. "We validated coordinated operation with nine physical Unmanned Aerial Vehicles (UAVs), extended the scalability study to thirty simulated UAVs, mapped 1.2 hectares in a single swarm flight, achieved an inter-to-total loop-closure ratio of 0.53 for collaborative Simultaneous Localisation and Mapping (SLAM), a formation similarity distance of 0.27, and approximately one-second inter-UAV command latency in the operational communication stack." The work reached Technological Readiness Level (TRL) 5 for collaborative multi-UAV mapping, validated beyond the laboratory, though not yet a finished commercial product.

Together, the five OpenSwarm proofs-of-concept – renewable-energy communities, forestry drones, marine noise monitoring, industrial health and safety, and communications aboard moving trains – tested whether one architecture could serve static and mobile nodes, indoor and outdoor settings, and markedly different operational demands. That breadth gives the consortium confidence that the approach generalises.

Application and implementation

OpenSwarm's reusable legacy is not a single product but a modular software ecosystem, combining open-source components, documented interfaces and partner-specific technologies, organising communications, AI, simulation, robotics and swarm technologies around a common architecture. It separates capabilities – including localisation, navigation, mapping, coverage and collective behaviour – from their enablers, and common code from platform-specific components. Because industrial users rarely adopt an entire research framework at once, components can be reused in stages: a manufacturer might begin with a networking component, low-power AI method, simulation tool or ROS integration pattern.

"The objective was not to force every platform to run exactly the same software," Watteyne explains. "It was to make



A DotBot demonstration at the Swarm Workshop in Brussels, 26–27 November 2025



Multi-UAV deployment in a forest setting: an OpenSwarm proof-of-concept demonstration of coordinated swarm formation

the relationship between technologies understandable: what a component does, what it requires, how it interfaces with other components and where it has already been validated." The DotBot testbed is also an open platform used by several research groups and institutions. Researchers and students can programme the robots remotely and run their own experiments.

Turning that foundation into industrial technology, Watteyne says, requires the "difficult details" to become predictable: stable Application Programming Interfaces (APIs) and hardware abstraction; cybersecurity across the device lifecycle; and continued standardisation, including through the LAKE (Lightweight Authenticated Key Exchange) working group at the IETF. Once a system runs unattended, reproducible configuration, containerisation, automated logging and remote diagnostics matter as much as any algorithm. Commercial integration remains sector-specific: forestry swarms must meet UAV regulation and fit Geographic Information System (GIS) workflows; marine systems need robust packaging and credible environmental data products; and factories face their own interfaces and certification requirements. "OpenSwarm provides the technological foundation," Watteyne says. "Industrialisation now requires the surrounding engineering ecosystem."

Couceiro sees the clearest near-term value in field robotics. "Agriculture, forestry, environmental monitoring, construction and aquatic operations all involve large physical spaces in which using one highly sophisticated robot is not always the optimal solution," he says. "Multiple smaller systems can provide spatial coverage, redundancy and parallelism, provided they can coordinate intelligently."

Lessons learned

A central lesson from the five proofs-of-concept was that the network, computation, energy and physical environment cannot be designed independently. A forest and a factory floor look nothing alike, yet both require systems to adapt as conditions change.

The industrial health-and-safety proof of concept showed why. Traditional safety systems, Watteyne notes, are usually built around fixed zones, predetermined alarms or centralised monitoring; a network of intelligent devices can instead build a more contextual picture of a site and adapt as conditions change. In a construction setting, he says, "the important question is not simply whether a sensor detects a person or a machine, but rather whether the system can understand the interaction between them, share that information with the relevant agents and adapt the operation safely."

For factories and logistics systems, Watteyne distils the project into five principles: process data locally rather than sending everything raw; design for degraded rather than perfect communications; separate SLAM, planning, formation, communication and control through clear interfaces; invest as much in deployment and observability as in algorithms; and leave autonomous agents to handle local complexity while people retain mission-level supervision. "Those principles," he says, "are as applicable to a factory fleet as they are to UAVs operating in a forest."

Looking ahead

For Watteyne, what lasting success means for the project is clear. "For me, it would mean that five years from now people are using pieces of OpenSwarm without necessarily thinking of them as research prototypes," he says. He envisages an industrial partner folding a networking technology into a product, a robotics company reusing swarm coordination or energy-aware methods, a university teaching distributed robotics with DotBots, or a new project building on the

architecture instead of starting from zero. Education, he argues, matters as much as commercialisation because swarm systems combine disciplines (robotics, embedded systems, networking, AI, control and cybersecurity) that are still too often taught in isolation. "Europe benefits," he concludes, "if it controls not just individual sensors or AI algorithms, but the architectures and skills required to make thousands of intelligent devices cooperate efficiently and sustainably."

For Ingeniarius, field deployment is the next step. Having reached TRL 5 with collaborative multi-UAV mapping, Couceiro believes the route forward runs "through pilot deployments and Mapping-as-a-Service" to gather operational evidence and mature the technology towards TRL 6 and TRL 7, extending beyond the original forestry use case into biodiversity monitoring, land management and potentially agriculture. The company's involvement in the marine proof of concept has also opened a further direction in distributed aquatic sensing, smart buoys and unmanned surface vehicles.

For Couceiro, that future depends less on any single technology than on making separate fields work together. "Networking, constrained AI, energy management and swarm behaviour were not treated as independent demonstrations; the project forced them to meet real systems," he says. It also points to the next task: not proving another algorithm in isolation, but "engineering complete, dependable and economically meaningful systems in which many intelligent agents can work together in the real world."

The broader European opportunity lies wherever many physical assets must cooperate but energy and connectivity cannot be assumed. "Europe is not going to obtain sustainable digitalisation simply by deploying more computation everywhere," Couceiro says. "These systems have to decide where computation is worth performing, what information is worth transmitting and how collective resources should be allocated."

The project's most important legacy may therefore be a change in what industry expects from connected devices. Instead of thousands of sensors reporting to a central server, there will be collectives able to interpret, negotiate and act – not as a crowd of identical machines, but as a resilient system whose intelligence grows through collaboration.



This forestry multi-UAV proof-of-concept project is designed to support human labourers in wildberry harvesting.

PROJECT INFORMATION

Project Title
OpenSwarm – Orchestration and Programming ENergy-aware and collaborative Swarms With AI-powered Reliable Methods

Project Objective
OpenSwarm aims to enable energy-aware swarms of collaborative smart nodes, benefiting the environment, industry, and society. The project advanced low-power wireless technology, networking and management of mobile nodes, swarm programming frameworks, and collaborative energy-aware AI to build efficient, data-centric, trustworthy systems. It produced an open-source code base and validated benefits through five real-world applications in renewable energy, environmental monitoring, industrial health, and mobility.

Project Duration and Timing
The project ran from January 2023 to April 2026.

Project Funding
HORIZON-CL4-2022-DATA-01-03
€5,322,116.25

Project Partners
Inria (France, coordinator), Analog Devices (Ireland), IMEC (Belgium), Ingeniarius (Portugal), KU Leuven (Belgium), Siemens (Germany), Siemens (Austria), Siemens (Portugal), the University of Sheffield (United Kingdom) and Wattson Elements (France).

MAIN CONTACT

Name
Micael Couceiro
micael@ingeniarius.pt

Thomas Watteyne
thomas.watteyne@inria.fr

Phone
+33 6 58 68 43 52

Web address
https://openswarm.eu/





OpenSwarm

This project has received funding from the European Union's Horizon Europe research and innovation program under the grant agreement **No 101093046**.

CONNECTED INTELLIGENCE

MYRTUS: Turning the computing continuum into a cognitive working infrastructure

Cyber-physical systems increasingly depend on computing resources distributed between devices, local infrastructure and the cloud. The Horizon Europe MYRTUS project is developing an open and adaptive environment capable of orchestrating those resources as a single computing continuum. The project's scientific coordinator **Francesca Palumbo** and exploitation manager **Alessandra Bagnato** explain how its reference infrastructure, MIRTO cognitive engine and design tools have turned an emerging concept into an integrated system, while exposing the technical questions that remain unanswered.

A road camera and a person completing rehabilitation exercises at home appear to have little in common. Yet both rely on services whose computing demand can change sharply. Rush-hour cameras may process far more information than at midnight, while a rehabilitation platform may have one patient connected at one moment and several at another. In both cases, resources must respond without interrupting the service.

Keeping processing close to the user reduces latency and can protect sensitive data, but edge devices have limited capacity. A remote cloud offers more power, but increases communication time and may raise concerns about privacy, cost and energy use. Fog computing adds an intermediate layer. MYRTUS is developing a continuum in which workloads can move across all three.

Its purpose is to give cyber-physical systems what the consortium calls a "living dimension": the ability to observe changing conditions, reason about them and adjust their use of computing resources while they continue operating.

Towards a computing continuum

MYRTUS is part of a wider European effort to shape the infrastructure for distributed digital systems. As more processing takes place inside factories, vehicles, medical devices and urban systems, interoperability and control over data and computing resources become industrial as well as technical concerns.

"Europe is investing a lot in open, interoperable and trustworthy digital infrastructures," says Francesca Palumbo, MYRTUS scientific coordinator and an associate professor of electronics in the Department of Electrical and Electronic Engineering at the University of Cagliari. "Europe is really strong on edge devices. We are less powerful on the cloud side, which is a domain dominated by other countries."

Palumbo sees the continuum as a way to connect Europe's edge expertise with fog and cloud capacity. Workloads could draw on specialised devices, nearby servers or larger cloud clusters, while open interfaces reduce dependence on one provider. The best place to execute a workload can change with demand, particular workload characteristics, network conditions or energy use.

Traditional deployment decisions are largely static: programmers decide where a service should run and give an orchestration platform instructions for maintaining that arrangement. MYRTUS instead creates a closed loop. Monitors



collect information about the application and infrastructure, feed it into a distributed knowledge base and allow the system to assess whether the deployment remains suitable.

"In the past, you were used to addressing these problems statically," Palumbo explains. "But for many different reasons, systems need to adapt. That's the 'living dimension', meaning that the system self-adjusts automatically."

The aim is to respond to actual conditions, perhaps keeping processing beside a camera for low latency, moving it when an edge device reaches capacity or selecting a more energy-efficient resource.

Making systems cooperate

The difficulty in getting systems on the three layers to work together like this begins with heterogeneity. Edge devices, fog nodes and cloud servers do not necessarily use the same processors, instruction sets, operating environments or management and monitoring tools. They may have been supplied by different vendors and installed at different stages of an organisation's development.

"When you put together a big server, local computing nodes and edge nodes with different architectures, instruction sets and vendors, making them talk is difficult," says Palumbo. "It is like putting together people from different countries. Either they speak the same language, or you need a translation layer between them."

Cloud and edge systems are also modelled differently. Kubernetes can deploy and maintain containerised services, but follows rules supplied by developers rather than deciding whether another deployment would better balance performance, energy, privacy and security.

Palumbo also points out that there is no single, widely accepted standard for describing and operating an entire continuum. "MYRTUS therefore had to create enough common structure for its components to cooperate without making the result dependent on one hardware platform, cloud provider or application," she adds.

A workload placed close to its data may offer low latency and stronger privacy, while a larger server elsewhere may complete it more efficiently. Energy consumption, network availability, security and the capabilities of existing equipment all affect the choice. As Palumbo puts it, "It is a matter of trade-offs. There is no free lunch."

A common language

MYRTUS approaches the problem through three connected elements: a reference infrastructure spanning edge, fog and cloud; the MIRTO cognitive engine that orchestrates it; and a design and programming environment through which applications and their requirements can be described and compiled to a variety of devices.

The infrastructure uses Kubernetes as a common software foundation. Liqo, an open-source technology developed by consortium partner ArubaKube, links separate Kubernetes clusters and allows them to share resources. Cloud, fog and

edge resources can consequently appear as parts of one virtual environment even when physically distant. "Liqo is capable of bridging these distances and allowing virtual access," Palumbo says. "It is as though the resources are in the same room, with everyone sitting at the same table."

"MIRTO adds the decision-making layer," she continues. "It receives information about available nodes, their capabilities and their current utilisation. It also considers the requirements attached to an application. Its managers can then choose where a workload should run and reconsider that decision when conditions change."

The engine uses several optimisation approaches, including swarm intelligence and federated learning. Its modular architecture allows different applications to prioritise latency, energy use or specialised hardware.

"We are using AI not for users, in the sense that we are not creating an AI service for the application," Palumbo explains. "We are developing an AI-based orchestrator that serves the needs of the infrastructure itself. It creates the glue logic and the continuous optimisation that allow the system to remain as close as possible to its optimum."

MIRTO orchestrator is designed as a composable engine to allow organisations to select the MIRTO managers most relevant to their circumstances or replace one with a component designed around another optimisation goal. Standardised interfaces allow components written in different languages to communicate through the same architecture.

Applications are described using TOSCA, the Topology and Orchestration Specification for Cloud Applications, which MYRTUS has extended towards fog and edge resources. A TOSCA model describes an application, its resource requirements and its constraints. Modelio-based tools create these descriptions and support threat modelling, code generation and deployment.

Alessandra Bagnato, a research scientist and head of Modelio Research at SOFTEAM Software, part of the Docaposte Group, is MYRTUS exploitation manager. She values TOSCA because its models are portable across tools. "The consortium has released reusable MYRTUS TOSCA nodes so that other computing-continuum projects can build on the same descriptions," she says.

"If all the projects can share their nodes, that can give a real push to the standardisation effort," she adds. "In our case, we are open source, so it was relatively easy to put this in place."

Building an infrastructure

At the time of the interview, MYRTUS had integrated computing resources provided by partners at sites across Europe. These included a cloud cluster at TNO and a micro data centre by





HIRO in the Netherlands; two nodes supplied by Abinsula Srl in Turin and Sardinia, Italy; and heterogeneous edge-computing platforms provided by the University of Cagliari in Sardinia and the Universidad Politécnica de Madrid in Spain.

At present, the physical research infrastructure is being used by project partners and the end users are involved in its assessment. The underlying software and methods are intended to help other organisations reproduce or adapt the approach. "We have a good level of integration across countries, platforms and layers," Palumbo says. "The infrastructure has been fully developed and integrated. Now assessment with the end users will be fundamental. Looking further ahead, we could also envisage the computing continuum being offered as a service."

MYRTUS has connected heterogeneous clusters and orchestrated resources across them. However, its two application scenarios are still undergoing final assessment, with evidence on availability, energy efficiency and other indicators yet to be completed.

The result is therefore an integrated technical foundation rather than a finished commercial platform. "When we started the project, the continuum was almost a keyword," Palumbo says. "Now it is something concrete. In three years, we went from a buzzword to an infrastructure."

Testing the continuum

The smart-mobility scenario uses connected cameras around road intersections to monitor traffic and recognise dangerous situations. Processing beside a camera reduces latency and keeps sensitive data local, but high demand may exceed its capacity. MYRTUS can move less sensitive parts of the workload elsewhere.

"The closer you are to the edge, the better," says Palumbo. "At the same time, that is not always possible. Some computation has to be shifted elsewhere."

The healthcare case supports cooperative telerehabilitation using virtual-reality headsets and sensor-equipped garments. An exercise can involve up to eight participants, with virtual agents filling empty places. People should be able to join, leave or lose their connection without ending the session, requiring real-time performance, availability and reconfiguration.

This scenario highlighted a question that remains open. Kubernetes technologies are well suited to moving stateless services, but an active rehabilitation session contains state: the

system must remember its participants, events and progress. Moving the computation would therefore require the system to preserve and transfer that context.

"The use cases assess our technologies and provide feedback on the questions that remain open," Palumbo says. "In virtual telerehabilitation, moving the state remains a challenge. How to preserve the context of a live session while moving the computation is a direction we are interested in exploring through future research."

MYRTUS is application-agnostic: its use cases test a general infrastructure rather than determine it. The developed technology could support factory cameras, robots and collaborative robots running machine-vision or speech workloads. Instead of replacing older devices with insufficient capacity, an orchestrator could distribute their work between the factory floor, local servers and the cloud.

"When you add new cameras to a factory, the older cameras may be less capable of handling the computation," Palumbo explains. "You can dispatch part of that computation elsewhere. The orchestrator is not bound to the particular cameras or FPGA platforms that we are using in the project."

Adopters would need to describe their application in TOSCA, operate compatible Kubernetes and Liqo infrastructure, and make monitoring data available to MIRTO. Much of the software will be open source, allowing users to adapt it to their equipment, while individual components can continue through existing communities and commercial partners.

Bagnato says the consortium began its exploitation work by asking what each partner wanted to achieve, whom it intended to serve, and how its contribution could continue after funding. This defined value propositions for individual technologies and for their combined operation.

"All the bricks have their value, and the fact that they can be used together is the real power of MYRTUS," she says. "We have followed this approach from the very beginning because we wanted the partners to focus on what they wanted to achieve, especially the uptake of the results in the future."

The project website describes the results, many of which are open source and available to download, and they provide routes to their developers. Open availability does not by itself create a sustainable model. "The developers are offering their products, but they are also providing consultancy, training and support to sustain the business model," Bagnato explains.

"This allows open-source components to support commercial services, while companies contributing proprietary technologies can use the project demonstrations as evidence that their products operate within a wider computing continuum," she concludes.

Integration lessons

The project's direction has remained consistent, but its priorities have changed. Initially, computing capacity, network performance, energy, latency, security and trust were broadly equal drivers. "We realised that we had to prioritise," Palumbo says. "At the moment, we are optimising computing capacity and energy more than network bandwidth, because our use cases are more oriented towards that kind of demonstration."

MIRTO's modularity still allows a future user to prioritise network optimisation. The lesson is that a general architecture needs clear priorities in each deployment; attempting to optimise everything equally may prevent a useful decision.

Palumbo believes that the stateful telerehabilitation problem delivered a second lesson: integration exposes difficulties that are invisible while components are developed separately. "Connecting clusters is different from moving a live service whose users must not notice that anything has changed," she says. "The use cases are consequently more than demonstrations. They are a means of finding the boundary between what the continuum can currently do and what still requires research."

MYRTUS grew from earlier European projects, including CPSwarm and CERBERO, which provided technical foundations and trusted relationships. "Trust is not important only for the continuum; it is also important for projects," Palumbo says. "MYRTUS depends on partners allowing technologies developed in different organisations to interact through shared models and interfaces." Its organisation mirrors the interoperability it is creating technically.

The next wave

The project's final stage is to complete application assessment, publish its open sourcce software and connect the results with communities that can maintain and extend them. Its contribution to standards may prove as important as any single demonstrator.

MYRTUS is working with Gaia-X on how nodes authenticate and receive credentials to join the continuum. Palumbo says discussions are under way about using the project's approach as a reference process, allowing MYRTUS to contribute to common practice rather than only adopt it.

Open source supports that strategy by making the system's operation visible. For Bagnato, transparency is particularly important when software makes decisions about workload placement. "With a tool that has the intelligence of MIRTO, it is important to know what is behind it," she says. "Open source is a way to say: trust us, we are transparent. Whatever is happening behind the scenes, you can see and check."

More autonomous, agentic AI could eventually assume a larger role, but the consortium does not treat it as an immediate

objective. MIRTO relies on defined models, managers and optimisation processes whose actions remain controlled and inspectable.

"At the moment there is no agentic system taking decisions independently," Palumbo adds. "There is a model, and everything is controllable. Agentic approaches could be a direction to investigate, with all the precautions that would require. What we have now is something stable and controlled that does its job well."

Bagnato sees that stability as a foundation for future AI services, whose infrastructure must allocate more computation without sacrificing trust, efficiency or control.

The consortium does not claim to have solved the continuum. Stateful migration remains difficult, interoperability is developing and widely accepted standards have yet to emerge.

Although several challenges remain, the project's industrial partners have identified promising real-world applications and will now work through their marketing departments to identify markets ready to adopt these technologies. "The continuum is still young and not completely standardised," Palumbo concludes. "There are still different challenges to be faced. But we like working together, so if the next opportunity is not already there, we will create one."

MYRTUS's most important achievement may be to have made that next stage possible. It has connected real resources across cloud, fog and edge; built an engine capable of reasoning about their use; and established open interfaces through which other technologies can participate – meaning the computing continuum is now a real infrastructure that can be tested, challenged and improved.



PROJECT INFORMATION

Project Title
MYRTUS – Multi-layer 360° dYnamic orchestration and interopeRable design environmenT for compute-continUum Systems

Project Objective
MYRTUS develops an interoperable computing continuum integrating federated edge–fog–cloud infrastructure, MIRTO AI-driven orchestration, and an open-source Design and Programming Environment. It enables efficient, secure, and adaptive deployment of applications across heterogeneous hardware, with validation in smart mobility and healthcare/virtual telerehabilitation.

Project Duration and Timing
1 January 2024 – 31 December 2026 (36 months).

Project Funding
European Union Horizon Europe, Research and Innovation Action (HORIZON-CL4-2023-DATA-01-04). Grant agreement No. 101135183; EU contribution: €5,597,915.18.

Project Partners
Abinsula Srl (coordinator); University of Cagliari (scientific coordinator); Softeam; TU Dresden; TNO; University of Sassari; Lakeside Labs; Universidad Politécnica de Madrid; HIRO MicroDataCenters; Forge Reply Srl; ArubaKube; Canon Research Centre France; Università della Svizzera Italiana; King's College London.

MAIN CONTACT

Name
Prof. Francesca Palumbo
francesca.palumbo@unica.it

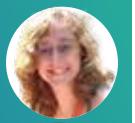
Alessandra Bagnato
alessandra.bagnato@docaposte.fr

Katiuscia Zedda
katiuscia.zedda@abinsula.com

Phone
+39 070 6755790

Web address
https://myrtus-project.eu/







MYRTUS
Multi-layer 360° dYnamic orchestration and interopeRable design environmenT for compute-continUum Systems

This project has received funding from the European Union's Horizon Europe research and innovation program under the grant agreement **No 101135183**.

CONNECTED INTELLIGENCE

HYPER-AI: Rethinking how cloud, edge and IoT work together

The cloud, edge devices and the Internet of Things contain enormous reserves of computing power, yet they are still commonly managed as separate layers. HYPER-AI is developing a platform in which heterogeneous resources can describe their capabilities, form temporary computing swarms and continuously adapt where data is processed. Part of the project's coordination team **Thanasis Moustakas** of **CERTH** explains how the project is combining semantic abstraction, decentralised AI and self-managing functions to build a more responsive computing continuum – and how manufacturing, mobility, healthcare, energy and agriculture will test this idea.

A cognitive computing continuum sounds abstract, but the idea behind it starts with familiar infrastructure. A central cloud can provide vast processing and storage capacity. Edge systems – a factory server, a vehicle computer or a local gateway – can act closer to where data is produced. Internet of Things devices add sensors, cameras, controllers and other constrained machines at the furthest reaches of the network. Together, these resources form a continuum from centralised computing to the physical world.

The problem is that the continuum rarely behaves as one, as its layers are typically managed by different systems, described in different ways and used through hierarchical decisions that begin in the cloud and are passed downwards. HYPER-AI – the Hyper-Distributed Artificial Intelligence Platform for Network Resources Automation and Management – is seeking to replace that fragmentation with a common, adaptive environment.

"The basic concept of the project is to make the computing continuum, which is comprised of different layers such as cloud, edge, Internet of Things (IoT) and even fog, behave like one intelligent pool of computing resources rather than as separate layers," explains Thanasis Moustakas, a research assistant at CERTH and member of the project's coordination team. "We try to unify behaviour across resources that could be anywhere from the cloud to IoT."

HYPER-AI represents each participating resource as a smart virtual computing node. The node describes what it can provide – computation, storage, connectivity or specialised capabilities – and can cooperate with other nodes to support an application. The 'cognitive' element lies in its capacity to use information about itself and its surroundings as the basis for action.

"A device starts by observing the environment, understanding it and, based on these assumptions, it decides and adapts," says Moustakas, offering a deliberately simple definition. "The objective is not to create one larger centralised system. It is to make intelligence and resource management available throughout the continuum, so the most appropriate combination of devices can be assembled and changed as conditions evolve."



HYPER-AI presentation at Swarm Intelligence Projects Workshop 2024 organised by the European Commission

Computing swarms

That change matters because sending every decision and every data-intensive task to a distant cloud is not always efficient. It can introduce delay, consume bandwidth, expose an application to connectivity problems and use energy moving data that might have been processed locally. A cloud may still be the right destination for a demanding workload; HYPER-AI's point is that it should not be the automatic starting place for everything.

"In traditional approaches, we make a decision in the cloud layer and then propagate decisions to the next layer," Moustakas says. "In our approach, we do it all at once; the decision-making is not hierarchical. We can move computation and decision-making to the most desirable place at any time. We do not necessarily have to start in the cloud."

A computing swarm is the temporary operational expression of that idea. HYPER-AI selects a group of nodes from across the available cloud, edge and IoT resources to perform a particular application or mission. The group is not a fixed cluster that must remain intact regardless of circumstances. It exists for the job, can be reconfigured while the job is running and dissolves when the work is complete. A node may even contribute to more than one swarm at the same time if its capacity allows.

"The group is something that is decided based on the goal at hand," Moustakas explains further. "It may be reconfigured as the dynamic nature of the setting triggers a need to reform it. Its lifetime should end when the application or mission ends, and then the nodes can go and form another one."

The difficult part is making those decisions while almost every relevant condition is moving. Devices vary in processing power, memory, storage, architecture, connectivity, ownership and availability. Applications bring requirements for latency, quality of service, data locality, security, energy and cost. At the same time, nodes can become overloaded, connections can deteriorate, and demand can change. Optimising one variable in isolation may simply move the problem elsewhere.

"The combination is the hardest part," says Moustakas. "The resources differ in characteristics such as compute capacity, storage, connectivity and availability. Applications also have requirements in terms of latency and quality of service. The main challenge is making good decisions when network conditions, resources and application requirements are changing at the same time. HYPER-AI therefore treats placement as a multi-objective optimisation problem rather than a single scheduling choice."

A self-managing continuum

The project builds on Kubernetes, the widely used system for orchestrating containerised applications, but extends it beyond the comparatively uniform servers for which such platforms were originally designed. A laptop or server capable of running Kubernetes can register as a native node. A smartphone, sensor or constrained IoT device cannot necessarily host the full software stack, so HYPER-AI has developed open connectors that allow it to be represented and managed within the cluster. Once registered, a device is described through a common semantic resource model. The same principle is applied to the application: instead of being written for one named machine or provider, it is represented through requirements such as processing, memory, storage, networking and accelerator needs. This creates a negotiable layer between what the application needs and what the distributed infrastructure can supply.

"The goal is to abstract things as much as possible," Moustakas says. "We have a common description for all devices with things like CPU, RAM and computing resources. For every application, we describe how much CPU is needed or whether an accelerator is required. The user does not need to understand every underlying device."

From there, HYPER-AI follows a continuous chain of observation and decision. State-estimation and anomaly-detection models give nodes visibility of their own condition and relevant parts of the wider system. A formation mechanism assembles an application-specific logical cluster. Within that cluster, nodes can bid for tasks, while optimisation models consider available resources, application constraints and data locality when placing the workload.

These decisions are supported by AI trained and tested initially in simulation. Moustakas developed a simulator able to generate nodes from different layers of the continuum and applications with different requirements, then used reinforcement learning to form logical swarms. The resulting policies can be compared with established baselines including round-robin allocation, first come first served, best fit and the default Kubernetes scheduler. This makes it possible to test behaviour repeatedly before placing the system in a live pilot, while also exposing the inevitable gap between a controlled model and an industrial network.

The project brings the management functions together through the self-CHOP principle: self-configuration, self-healing, self-optimisation and self-protection. Self-configuration selects and configures resources. Self-healing responds to unavailable or failed nodes. Self-optimisation improves resource use and workload placement. Self-protection detects threats and maintains trust. These are not four independent tools used at different points in a deployment; they form a feedback loop.

"This is a continuous procedure, and one thing can help the other," Moustakas explains. A missing or overloaded node is treated as a normal event rather than an exceptional failure.

"It is not a once-and-done procedure," he continues. "If the model detects that a node carrying a task has gone down, it must find another node to include in the swarm. If that is not feasible, it can reconfigure the whole swarm from the beginning."

Security must be distributed just as the computation is. HYPER-AI combines identity and authorisation, encryption and privacy mechanisms, distributed-ledger-based trust and AI-enabled intrusion detection. The design constraint is significant: a security



mechanism that works only on a server cannot protect a swarm containing constrained edge and IoT nodes. "The mechanisms have to fit the capabilities of all nodes in the system," Moustakas says. "The models are small enough to run even on small, constrained devices."

One architecture, five demanding tests

A platform designed to abstract heterogeneous resources could appear convincing in a single controlled setting and yet still fail when exposed to a different sector. HYPER-AI is therefore preparing five use cases in manufacturing, farming and agriculture, green energy, healthcare, and mobility and automotive. They have been selected as they do not share the same devices, data flows, urgency or performance measures, so transferability can be tested.

"The general idea is generalisability," Moustakas explains. "By using five distinct use cases, we aim to test whether this unified architecture can accommodate very different heterogeneous resources and very different application descriptions and needs.

"Agriculture includes image-based analysis for crop health and precision farming," he continues. "Healthcare focuses on early disease diagnosis and control. The green-energy case considers efficient data processing for monitoring critical infrastructure. In mobility, a connected vehicle's unused computing capacity can potentially support workloads while the car is parked or charging rather than driving."

For manufacturing, the continuum already exists in physical form in many industries on the factory floor. A modern plant may combine robots, machine controllers, sensors, cameras, augmented-reality glasses, local servers and cloud services. Waiting for a remote cloud is not necessarily appropriate when an operator needs immediate guidance or a robot requires a time-sensitive response. Local processing can reduce latency and dependence on connectivity, while the cloud remains available for tasks that genuinely need its scale.

"The HYPER-AI manufacturing case combines AR- and AI-assisted remote assembly with Robot-as-a-Service," says Moustakas. "A worker wearing AR glasses can receive instructions from a remote expert, while robots and local infrastructure gain access to a wider, coordinated pool of computing resources. Crucially, the glasses are not treated merely as a display at the end of the system. They become a device node that can be connected through the project's open connector and participate in the wider cluster.

"In a traditional setting, that device would not have a way to connect to the Kubernetes cluster," he continues. "We provide a way to connect it so it can work with the other devices already in the cluster and make decisions that benefit the whole system."

The project is not redesigning the worker's application or deciding what instruction should appear in the glasses; its role is to match the application's requirements to the resources capable of delivering it. That boundary matters when assessing implementation: HYPER-AI addresses computing orchestration, while training, interface design and process change remain responsibilities for the industrial deployment.

An integrated platform

At this stage in the HYPER-AI's lifespan, things have moved

well beyond architectural intent. Public outputs now include its distributed adaptive cloud–continuum architecture, semantic resource and application models, the Hypertool abstraction and lifecycle–management framework, learning–based scheduling research, a Kubernetes workload simulator and LCF–Bench for evaluating logical–cluster formation. Individual technical components have stable versions and have been evaluated separately. The decisive next step is integration.

Moustakas is careful not to present simulated component performance as proof of an end–to–end platform. ''We have a stable version for each tool, but as integration continues, we have not yet tested them all together,'' he says. ''When designing this, we made decisions about how inputs would proceed from one tool to another. The thing that remains is to see how the system achieves results in a real setting and how one model is affected by another.''

That distinction is central to the project's outcome. HYPER–AI aims to improve measures such as latency, task completion, resource utilisation, energy use, financial cost, resilience and security, while making more effective use of edge nodes. At the interview stage, however, those project–level KPIs had not yet been measured in the pilots. ''We have simulation results for each model separately,'' Moustakas says. ''The KPIs are not being measured yet because we have not started pilot testing.''

Integration is more than a software assembly exercise. A scheduling policy that performs well alone may behave differently when it receives imperfect state information, competes with security checks or triggers a swarm reformation. A connector that registers a device must remain reliable as the network changes. The pilots will therefore test not simply whether every component runs, but whether the closed loop remains stable and useful in environments it cannot fully control. This is where HYPER–AI's central proposition — cognition across the continuum — must become measurable operational behaviour.

Implementation

HYPER–AI's common descriptions and open connectors are intended to lower the effort needed to introduce a new resource or application. For the broad range of devices already considered by the project, Moustakas does not expect adoption to depend on a wholesale hardware replacement. ''An organisation still needs suitable resources and a working application, and that application must be represented in a form the HYPER–AI mechanisms can understand,'' he says. ''The platform then provides the abstraction and coordination layer between them.''

The code is being made available under the Apache License 2.0 through the Eclipse Foundation's GitLab environment. ''Basically, anyone can have access to our code,'' Moustakas continues, while noting that individual partners retain rights associated with their pre–existing and separately developed assets. ''Open source makes inspection, reuse and further development possible; it does not by itself supply deployment engineering, support, governance or a commercial service model.''

That is why the most credible route into operation may begin with parts of the platform rather than a single all–encompassing product. ''There is commercial potential, maybe not as a whole, but we could start with one mechanism or a combination of mechanisms and package them in one release,'' Moustakas suggests. ''There is potential, but there needs to be work on that.''

Hypertool, open connectors, cluster–formation methods, scheduling or self–protection functions could each solve a narrower problem while the complete continuum matures.

For an industrial adopter, the business case will ultimately depend on the pilots. Better use of existing edge resources, reduced reliance on remote processing, faster response and greater resilience are plausible benefits, particularly in manufacturing. HYPER–AI's discipline is to convert those promises into evidence before claiming them as results.



HYPER–AI consortium meeting in Brussels

A moving target

The project's clearest lesson is that abstraction is necessary but insufficient. Giving every resource a common description does not remove changing availability, incomplete information or the interactions between discovery, lifecycle management, scheduling, security and recovery. "The computing continuum is much more dynamic and heterogeneous than we anticipated," Moustakas says. "Abstraction is not enough on its own. Resource discovery, lifecycle management, dynamic swarming and all the models we develop have to work together and respond to changing conditions."

One practical example came from choosing Kubernetes as the orchestration foundation. The early concept assumed that devices could be abstracted through one route. In implementation, constrained devices could not run the necessary Kubernetes software and therefore could not register directly. The answer was not to exclude them, but to create open connectors that mediate between the device and the cluster. The solution embodies the project's wider method: preserve a common operational model without pretending that every physical resource is the same.

The future of the continuum

The immediate programme is clear: bring the tools together in the project's physical cluster, connect the pilot devices and execute the five use cases. That work will determine whether the architecture can move from stable individual mechanisms to an integrated demonstrator, and whether optimisation objectives that look promising in simulation translate into worthwhile improvements in real settings.

Beyond the pilots, HYPER-AI's prospects depend on a wider European ecosystem. Moustakas identifies "open interfaces, interoperability and common representations" as the essentials. The project is working with related European initiatives and standardisation activities so that its resources and applications are not described in a private language that only one platform understands. Europe's opportunity lies not simply in building another orchestration stack, but in creating an open environment in which devices, providers and applications can cooperate without surrendering control to one central operator.

Success may therefore take more than one form and the complete HYPER-AI platform could be adopted in later projects or industrial environments. Individual components could be maintained, combined with other systems or developed into services. Research teams could use the project's models and benchmarks to improve their own approaches. "We would definitely like to see adoption of the whole platform or some components after the project ends," Moustakas says. "Another contribution would be for us, or someone else, to build on the existing components and evolve or improve them."



PROJECT INFORMATION

Project Title
HYPER-AI – Hyper-Distributed Artificial Intelligence Platform for Network Resources Automation and Management Towards More Efficient Data Processing Applications

Project Objective
HYPER-AI develops an open, secure, and interoperable AI-driven computing continuum architecture that enables autonomous edge intelligence, cognitive orchestration, and dynamic resource management across cloud–edge environments.

Project Duration and Timing
1 April 2024 – 31 March 2027

Project Funding
HORIZON-RIA – HORIZON Research and Innovation Actions
€ 4 628 975,00

Project Partners

EU Partners:
CERTH [EL] / ETHNIKO KENTRO EREVNAS KAI TEXNOLOGIKIS ANAPTYXIS
TID [ES] / TELEFONICA INVESTIGACION Y DESARROLLO SA
ECL [DE] / ECLIPSE FOUNDATION EUROPE GMBH
NKUA [EL] / ETHNIKO KAI KAPODISTRIAKO PANEPISTIMIO ATHINON
VIF [AT] / VIRTUAL VEHICLE RESEARCH GMBH
ODINS [ES] / ODIN SOLUTIONS SOCIEDAD LIMITADA
SABO [EL] / SABO SA
TRUST-IT [IT] / TRUST-IT SRL*
*8.1 COMMPLA SRL (AFFILIATED)
CUT [CY] / TECHNOLOGIKO PANEPISTIMIO KYPROU
EBOS [CY] / EBOS TECHNOLOGIES LIMITED
ENEA [IT] / AGENZIA NAZIONALE PER LE NUOVE TECNOLOGIE, L'ENERGIA E LO SVILUPPO ECONOMICO SOSTENIBILE

Associated partners:
CSEM [CH] / CSEM CENTRE SUISSE D'ELECTRONIQUE ET DE MICROTECHNIQUE SA – RECHERCHE ET DEVELOPPEMENT
HES-SO [CH] / HAUTE ECOLE SPECIALISEE DE SUISSE OCCIDENTALE

International partners:
SUNDO [S.KR] / SUNDOSOFT LTD

MAIN CONTACT

Name
Thanasis Moustakas (right)
atmoustakas@iti.gr

Niccolò Zazzeri
n.zazzeri@trust-itservices.com

Veronica Vuotto
v.vuotto@trust-itservices.com



Web address
https://hyper-ai-project.eu/

HYPER-AI
Revolutionising big data applications with autonomous cloud-to-edge resources

This project has received funding from the European Union's Horizon Europe research and innovation program under the grant agreement **No 101135982**.

ZERO DEFECT MANUFACTURING

FLASH-COMP Industry Event 2026

Executive Summary

As the FLASH-COMP project reached its conclusion in September 2026, it hosted a special industry event at the Avigliana manufacturing plant of project partner AZIMUT, the world-renowned yacht builder, near Turin, Italy. The event brought research and industry together to hear how the project had successfully demonstrated Zero Defect Manufacturing (ZDM) technologies on the factory floor, see those technologies in action and address a crucial question: what does it take to move a technically successful R&D result into reliable industrial use?

The following pages summarise the presentations and candid discussions that took place. They conclude by outlining a possible route towards the wider deployment of results such as those produced by FLASH-COMP on factory floors across Europe and beyond.



This project has received funding from the European Union's Horizon Europe research and innovation program under the grant agreement No. 101058458.

Chapter 1

How FLASH-COMP moves towards ZDM

Across marine and aerospace use cases, FLASH-COMP has developed and industrially validated an integrated approach to Zero Defect Manufacturing for liquid resin infusion. At Azimut, inspection tools, data infrastructure, real-time simulation and an operator-facing decision support system has been tested at industrial scale across the process sequence. The system checked the dry layup, monitored pressure and resin movement during infusion, and used hyperspectral imaging to follow gelation and cure. The platform brought this information together so operators could see conditions, receive alerts and consider action while there was still time to influence the outcome.

During the event, FLASH-COMP partners made it clear what this impressive achievement does and does not mean. The industrial campaign established feasibility and demonstrated an end-to-end technical chain in a real production environment. It did not turn every project result into a finished commercial product, prove repeatability across every component or production line, or settle the business and governance arrangements needed for joint exploitation. Several technologies reached reported readiness levels of six or seven, yet further work remains on robustness, cost, installation, data rights, integration, customer value and support. The present AZIMUT rig can address only part of the shipyard's production range. Larger and more complex moulds will need different equipment and further validation.

This distinction shaped the industry roundtable. Panellists argued that funded projects are often expected to reach industrialisation within budgets and timescales that can realistically deliver applied research and relevant environment validation. Projects may begin with a research concept and recruit end users later, when a stronger route would start



An AZIMUT yacht manufactured using a liquid resin infusion process close to being finished at the Avigliana factory.

with the manufacturer's problem and a credible post-project plan. Demonstrations also tend to use the best line and most motivated staff. Replication on older brownfield lines is a separate challenge.

The panel's message was constructive and articulated one of the key outcomes of the FLASH-COMP project, which has now ended. Manufacturers should understand their process, collect useful data, involve operators and define the cost of the defects that matter most. Technology providers should progress from description and diagnosis towards prediction and prescription in manageable stages. Research programmes should use realistic readiness claims and support the *costly work after a pilot*. Result owners should develop modular offers, quantify customer value and use further trials to build investment evidence.

The project's exploitation portfolio comprises five Key Exploitable Results, five further Exploitable Results and three Usable Results. The immediate task at the project's end is to decide which warrant further investment, agree how joint results will be owned and delivered, and turn targeted engagement into pilots, licensing discussions, validation partnerships, training use and research collaboration.

The event closed the loop from research to industrial demonstration while defining the harder work required for dependable use. Its final message was clear: **Zero Defect Manufacturing can support sustainability and European competitiveness together when technology, people, economics and implementation are treated as one industrial problem.**

Context and purpose

FLASH-COMP was established four years ago to make liquid resin infusion more reliable, resource efficient and human centred. Dry fibres and core materials are arranged in a mould, sealed under a vacuum bag and infused with resin. The process produces large, lightweight structures, but preparation remains largely manual and important quality information often arrives only after infusion. By then, a defect may require expensive repair or rejection.

The project's response has been to move quality control into the process. Monitoring instruments capture the state of the layup, vacuum, flow front, gelation and cure. A shared data architecture transfers and stores the information. Simulation and analytical tools compare expected and observed behaviour. A decision support system presents the findings to operators through a human-machine interface. The operator retains authority, but works with better evidence and earlier warning.

European manufacturers face pressure to improve productivity while reducing waste, energy use and environmental harm. A failed large composite component can consume substantial material, labour and production capacity. The barriers also recur across manufacturing: variable processes, fragmented data, legacy equipment, scarce skills, disruption risk and uncertainty over return on investment.

Holding the event inside AZIMUT's shipyard gave these issues industrial weight. Delegates saw the working environment, toured production and inspected the FLASH-COMP rig. The programme connected the research history with industrial evidence, supporting methods and the decisions that follow a pilot. Delegates who attended online were able to see the whole process on a film created by Insight Media at the factory, while material from the event has also been developed for Projects Magazine and EFFRA's Manufacturing Partnership Days.

Opening words

Carlo Ighina, head of R&D at AZIMUT, opened proceedings by welcoming delegates to AZIMUT and placed the event within the practical setting of yacht design and manufacture. AZIMUT produces vessels across a broad size range and works with complex geometries, high finish expectations and multiple materials. This setting is demanding for process innovation. Production volumes are significant for yacht building, but do not resemble the repetitive conditions of automotive mass production. Solutions must therefore cope with variety, manual work and the commercial importance of visible quality.

Project coordinator **Janusz Poplawski** of LORTEK then introduced FLASH-COMP and its aim of manufacturing composite parts right first time. The four-year Horizon Europe project began in October 2022 with a budget of about EUR 6.69 million and EUR 5.61 million in funding. Its two main industrial pathways addressed yacht components at AZIMUT and an aerospace wing box section. The opening project film explained the central idea: combine advanced monitoring with artificial intelligence and decision support so that a process problem can be identified and addressed before it becomes a finished part defect. The environmental ambition was equally prominent because avoiding scrap and repair also avoids unnecessary resin, consumables, energy and labour.

Ennio Chiatante of COMAU set the work within the European research and innovation system. Speaking as an EFFRA board member as well as from COMAU, he described EFFRA's role on the private side of the European manufacturing partnership and the substantial share of partnership funding that reaches industry as well as research and technology organisations. His intervention linked the project to a larger competitiveness question. Europe needs manufacturing capability, and continued public support depends in part on industry demonstrating the value of collaborative research. That theme returned later in the roundtable, where the panel examined whether present funding instruments provide enough support for the stages between a validated prototype and an industrial offer.

Chapter 2

The R&D that has delivered a platform for decision support

From the laboratory to the factory floor

Naval composite applications

Ricardo Jorge Rocha of INEGI explained how the marine pathway developed from a basic industrial problem. In conventional liquid resin infusion, quality checks often take place after defects have formed. Manufacturers protect themselves by using excess resin, carrying out extensive inspection and accepting repair as a routine part of production. The result is avoidable waste and late discovery of problems.

FLASH-COMP created a digital chain connecting sensors, a shared data space, a hybrid twin, decision support and the operator. Development progressed through a robot-assisted laboratory cell, a simpler repeatable rig and a pre-industrial environment with multiple geometries and more realistic lighting and process conditions. The digital testbed demonstrated near real-time acquisition and analysis for flow tracking, dry spot prediction, vacuum leak detection and defect severity assessment. Transferring this capability to Azimut required compromises that are instructive for industrial uptake. A hull was the first candidate, but workers routinely enter and move around the mould. That created unacceptable risks to sensors and to the consistency of the measurements required for validation. The consortium therefore selected a smaller hard-top cover of about six square metres. The choice did not diminish the technical achievement, but it defined its scope: the team proved the integrated concept on an industrial component under controlled conditions, while leaving hull and large deck deployment for a later stage.

The naval demonstration combined LiDAR and a profilometer before infusion, pressure sensors and matrix cameras during infusion, and hyperspectral imaging after infusion. Rocha concluded that the question had moved from whether the approach could work to how it could be scaled and implemented. The current rig was estimated to be compatible with roughly 20 to 30 per cent of Azimut's production line. Wider deployment will require larger equipment, additional sensing, shop floor integration, standardisation and a business case that compares the value of avoided waste and repair with the cost of hardware, software and support.

Aviation validation

Cristian Builes of AIMEN presented the aerospace use case at IAI, which examined a carbon fibre and epoxy wing box section. The work focused on fibre optic sensing and 3D profiling under the stricter structural and quality requirements of aviation manufacturing. Fibre Bragg gratings provided measurements at selected points, while distributed fibre sensing produced information along a continuous length. The development team had to ensure that sensing did not compromise the final laminate and that the installation approach remained compatible with manufacturing requirements.

The validation progressed in stages. Laboratory trials established what the sensors could detect during vacuum application, resin arrival, infusion and temperature change. A smaller wing box section was then used to refine the preform, consumables, inlet and outlet strategy and sensor positions. Full-scale trials on a 1.2 by 0.7 metre demonstrator exposed effects that had not appeared at smaller scale, including differences in vacuum distribution and dry areas. The team adjusted the setup and used sensor evidence to reach a good quality component.

The talk showed why scaling is a technical activity rather than a simple enlargement of a laboratory setup. Sensor signals require processing and interpretation, geometry must be linked to the data, and a monitoring arrangement can itself affect the process if it is poorly installed. By combining fibre optic data with a geometric point cloud and the digital platform, the aerospace work established a route from raw measurement to a process record that can support analysis, digital twins and future decision support.

Industrial implementation at AZIMUT

Matteo Lasagna, from AZIMUT, described the problem from inside the shipyard. AZIMUT introduced its vacuum-assisted resin infusion process in the early 2000s and has accumulated extensive practical knowledge across components ranging from small hatches to hulls, decks and superstructures. That knowledge is valuable, but much of the current monitoring remains offline. Operators use visual inspection, checklists

and basic instruments to examine fibre and core placement, temperature, humidity, vacuum, leaks, flow and cure.

FLASH-COMP introduced two complementary capabilities. Simulation can help design the infusion layout and predict resin velocity, pressure, filling time, temperature, cure and deformation. In-line monitoring can then compare the real process with that expectation and identify divergence early. AZIMUT adopted a staged pilot to validate the technology without disrupting production. Operators, laboratory technicians and process engineers contributed feedback on features and usability, and a two-day training course connected composite theory, Azimut's process and the complete platform. The resulting human-centred interface gathers outputs from the monitoring tools in one place and generates an alert when a parameter moves outside its defined threshold. Lasagna stressed that the system is intended to support the operator rather than replace practical expertise. Its prospective industrial value includes earlier quality control, less repair and waste, better working conditions and a traceable data record for every infusion. That record can also improve future process design, turning experience from one campaign into evidence for the next.

Readiness for adoption

Marta Pinzone of Politecnico di Milano broadened the discussion beyond technology. FLASH-COMP has developed a practical set of methods to help manufacturers assess whether they are ready for Zero Defect Manufacturing, identify where improvement matters most and anticipate the effect on work.

The Zero Defect Manufacturing maturity model assesses six dimensions: Product, Process, Platform, People, Partnership and Performance – the 6Ps approach. It allows an organisation to compare its current profile with a desired future state and prioritise feasible actions. The model matters because a sensor or platform cannot deliver value if the organisation lacks compatible processes, skills, data or ownership of the change.



The complex hull of a yacht on the factory floor where the fit-out takes place

The eco-efficiency method maps process time, labour, material cost and environmental impact. In the Azimut baseline presented at the event, post-infusion work accounted for about 57 per cent of total cost because of long, labour intensive operations. Resin infusion created the highest environmental impact, about 79 per cent of the total, mainly because of resin consumption. Vacuum bagging accounted for a further 13 per cent. These findings indicate where defect prevention and better control can affect both cost and environmental performance, while also showing that the benefits must be measured at process level rather than assumed.

The job, task and skills analysis examined which activities would remain, change or disappear when digital monitoring and decision support are introduced. Around half of the tasks reviewed were expected to change in some form. Open training resources were created in response. Together, the maturity model, eco-efficiency assessment and skills work turn supporting activities into part of the implementation method. They help a manufacturer decide what to change, how to measure the result and how to prepare the people who will use it.

A complementary perspective from TURBO

Nikos Pantelelis of Synthesites presented TURBO, a sister project of FLASH-COMP, which is addressing the infusion and inspection of offshore wind turbine blades. Its use case involved a proprietary one-shot process in which a complete blade, potentially more than 100 metres long, is infused rather than two halves being manufactured and bonded. Multiple resin fronts must reach and merge at the right locations, internal areas are difficult to observe, and the cycle must fit daily production.

TURBO combined wired and wireless sensors, flow simulation, machine learning, a digital twin and post-demould inspection. The final trials used a representative blade section rather than a complete blade. Pantelelis was candid that the project did not achieve every original ambition. Simulation work took longer than expected, reducing the time available to build a robust machine learning model.

Even so, the project produced useful sensing, control and inspection results and completed industrial demonstrations with Siemens Gamesa. The comparison with FLASH-COMP reinforced three points. Composite sectors require different sensing strategies and levels of assurance. A zero waste or zero defect ambition is a direction of travel rather than a claim that all variability can immediately be removed. Finally, individual project components may retain industrial value even when the complete proposed system needs further work.

Chapter 3

Validated technology delivering real-time intelligence

Inside the AZIMUT validation

The validation film showed how the components worked together. Following pre-industrial testing at INEGI, partners spent three days at Azimut preparing and monitoring a hard-top cover. A portable rig covered an area of approximately 5.4 by 5 metres, and motorised linear units moved instruments with a narrow field of view across the working area.

Before infusion, LiDAR and profilometer scans recorded the dry geometry and checked the layup under vacuum. Automated analysis identified pipes, folds, wrinkles, sensors and other obstructions, then proposed regions where cameras could see the advancing resin. A monitoring system cannot produce trustworthy information if its view is blocked or its installation creates a leak.

During infusion, pressure sensors and matrix cameras sent data through an edge unit to the data space. The platform displayed pressure, detected resin arrival and tracked the flow front. Simulation provided a short look ahead and highlighted possible dry areas. Operators compared the detected outline with the physical process and decided whether intervention was required. The platform did not yet close the loop through automatic valve actuation.

After impregnation, hyperspectral scans monitored chemical changes associated with curing. Repeated measurements over a selected line created a spatial and temporal view of cure and could support alerts for slow, uneven, insufficient or excessive curing. The complete campaign therefore demonstrated the project's intended information chain: inspect the prepared part, observe the live process, compare it with prediction, support an operator decision and retain a traceable record.

The film showed integration under factory conditions and exposed the work behind it. Equipment had to be positioned, connected and calibrated around normal materials and movements. Lighting, reflectivity, condensation, temperature and obstructions could affect measurement quality. The achievement was a functioning system in a defined context. Larger moulds, other resins and less favourable conditions remain the next test.

Production readiness

Jose Pedro Ferreira of INEGI explained the physical monitoring sequence in more detail. The rig provided a stable frame for LiDAR, the profilometer, cameras and process sensors, while allowing scanning movement where required. The team used



The hyperspectral camera and profilometer mounted on the FLASH-COMP rig placed over the hard top being infused



The hard top being infused, showing the sensors placed on the top of the blue vacuum bag before infusion takes place

two pre-infusion scans: one of the dry laminate and another after the vacuum bag, tubes and sensors were installed. Comparing them allowed the algorithms to distinguish the part geometry from later obstructions and select appropriate regions of interest.

Pressure sensors detected local vacuum conditions and resin arrival. The team designed mounting pieces that could sit above the vacuum bag while limiting the risk of introducing a leak. Matrix cameras captured the full working area and processed images in real time to trace the resin front. Pressure and flow information then supplied the simulation and decision support functions. Ferreira's presentation showed that production readiness includes the readiness of the measurement system itself: geometry, visibility, sensor placement, sealing and data transmission must all be verified before the first resin enters the mould.

Real-time decision support

Gorka Alvarez of MSI described the digital platform that connects shop floor data with analysis and the operator interface. A local edge machine receives information from cameras and sensors, performs processing where necessary and communicates with the FLASH-COMP data space using MQTT. The data space cleans and stores the information and makes it available to simulation, defect estimation and decision support tools. A web interface presents the result.

The platform organises digitalisation in four stages. Descriptive functions show what is happening. Diagnostic functions help explain an abnormal condition. Predictive functions estimate what could happen next, including flow and dry spot development. Prescriptive functions recommend what the operator should do. A fifth future step would connect the recommendation to actuators. At the event, Alvarez replayed an AZIMUT campaign, showing recorded alerts, pressure anomalies, infusion progress and post-infusion information. Historical replay and operator notes give the platform value beyond live control by supporting training, investigation and continuous improvement.

Post infusion quality

Project coordinator Janusz Poplawski of Lortek closed the technical sequence with hyperspectral cure monitoring. Unlike a conventional RGB camera, the selected near infrared system captures more than one hundred spectral bands between approximately 900 and 1700 nanometres. Changes in chemical bonds during curing alter the spectral response. Regression models can therefore relate the image data to degree of cure without contacting or damaging the part.

Laboratory work on aerospace resin showed good correlation with reference measurements and led to two scientific publications. At Azimut, the camera and lighting were mounted on the rig and used for periodic scans after impregnation. The results could be displayed through the decision support platform and used to identify local or systematic anomalies.

Poplawski also set out the commercial limitations. Models require broader evidence across resin mixtures, temperature, humidity and other conditions. The current analysis is largely superficial, while deeper measurement could extend its value. The linear camera is expensive and requires mechanical scanning. A dedicated matrix sensor covering only the most informative spectral bands could reduce cost and simplify installation. Applications involving ovens or autoclaves may offer a stronger energy case, but would also require protection from high temperatures. The presentation was a useful example of disciplined exploitation thinking: the scientific result is promising, but market adoption depends on a cheaper, more robust and application specific solution.

Chapter 4

Expert views on the way forward for ZDM projects

Industry roundtable on competitiveness and adoption

The final roundtable brought together Ennio Chiatante, Nikos Pantelelis, Peter Craamer of MSI and Carlo Ighina, with William Davis of Insight Media as moderator. It asked whether Zero Defect Manufacturing can strengthen European competitiveness while supporting sustainability. The discussion ranged beyond the project and became a frank examination of why research results so often stop after a successful pilot.

The manufacturing problem is specific

The panel first challenged the idea that there is one composite The panel first challenged the idea that there is one composite manufacturing problem. In yacht building, surface finish and rework may dominate the business case. In aerospace, structural assurance has a different priority. Wind blades, automotive components and closed mould processes introduce other geometries, access constraints, cycle times and levels of automation. The same underlying technology may therefore need different sensors, thresholds, evidence and commercial packaging.

Ighina described yacht composite production as an industry moving from artisanal practice towards industrial control. AZIMUT has decades of experience, but process knowledge still depends heavily on people and is not always captured in design tools or data. The sector combines complex products and meaningful production volume without the repeatability of mass manufacturing. There is no simple model from another industry that can be copied. Simulation, monitoring and process records are important because they can turn accumulated experience into a more predictable production system.

Useful data and human expertise

Craamer identified reliable data as the foundation of repeatability. **Pantelelis** added that the data must also be useful. Collecting large volumes without a connection to defects, decisions or customer value creates burden rather than insight. The panel called for process data, environmental context, supplier information and operator knowledge to be considered together. It also warned that departments often collect different measures for different purposes. A maintenance team, production manager and finance director may all describe the same issue through incompatible indicators.



One of the pressure sensors (left) on the infused part and the lidar camera (right) mounted on the rig, ready to monitor the part after the infusion has taken place



The practical sequence should begin with descriptive and diagnostic understanding. Manufacturers need to know what happened and why before expecting artificial intelligence to predict an outcome or prescribe an action. Most industrial work still lies in these first stages. Claims of predictive control can create unrealistic expectations if the underlying process is poorly understood or the data are inconsistent.

Human involvement was discussed from two directions. Operators hold experience that a digital system must capture and respect. Involving them in the design also builds trust and helps correlate data with factory reality. At the same time, a commercial system should not impose excessive new manual work or depend on one expert to keep it running. The goal is to retain human judgement while automating repeatable acquisition and presenting only information that helps a decision. Training and change management are therefore part of the technical deployment.

Sustainability and competitiveness

The panel did not assume that sustainability automatically creates competitiveness. **Chiatante** noted that European manufacturers cannot rely on sustainability alone when global customers also compare price and performance. Europe is unlikely to win every market through mass production volume, so quality, specialist capability and high value manufacturing become important differentiators.

In the processes discussed, however, the link can be concrete. Earlier defect detection avoids adding labour, machine time and energy to a part that will later fail inspection. Less repair means lower use of resin, fillers, paint and other materials, as well as less dust and fewer emissions from finishing work. Better cure control can reduce unnecessary heating in processes that use ovens or autoclaves. Shorter or more reliable cycles also release capacity. **Pantelelis** described an industrial wind blade application in which monitoring reduced cure cycle time, creating both energy and production benefits.

The investment case must consider the whole product and process rather than the budget of one department. A monitoring system may affect research and development, production, quality, sustainability and after-sales risk. If each function evaluates only its own cost, the investment may appear weaker than its company-wide effect. The converse is also true: broad claims about waste or productivity do not constitute a business case. A manufacturer needs an agreed baseline, implementation cost, realistic avoided losses and evidence that the system can achieve those benefits in its own environment.

The difficult path from pilot to market

The most hard-hitting exchange concerned research structure, funding and time. **Chiatante** acknowledged that the proportion of publicly funded results reaching industrial



Carlo Ighina delivers a talk about how the large hulls of AZIMUT yachts are infused with liquid resin delivered through a complex arrangement of pipes

Project Partners


Carlo Ighina


Cristian Builes


Janusz Poplawski


Ricardo Jorge Rocha


Ennio Chiatante

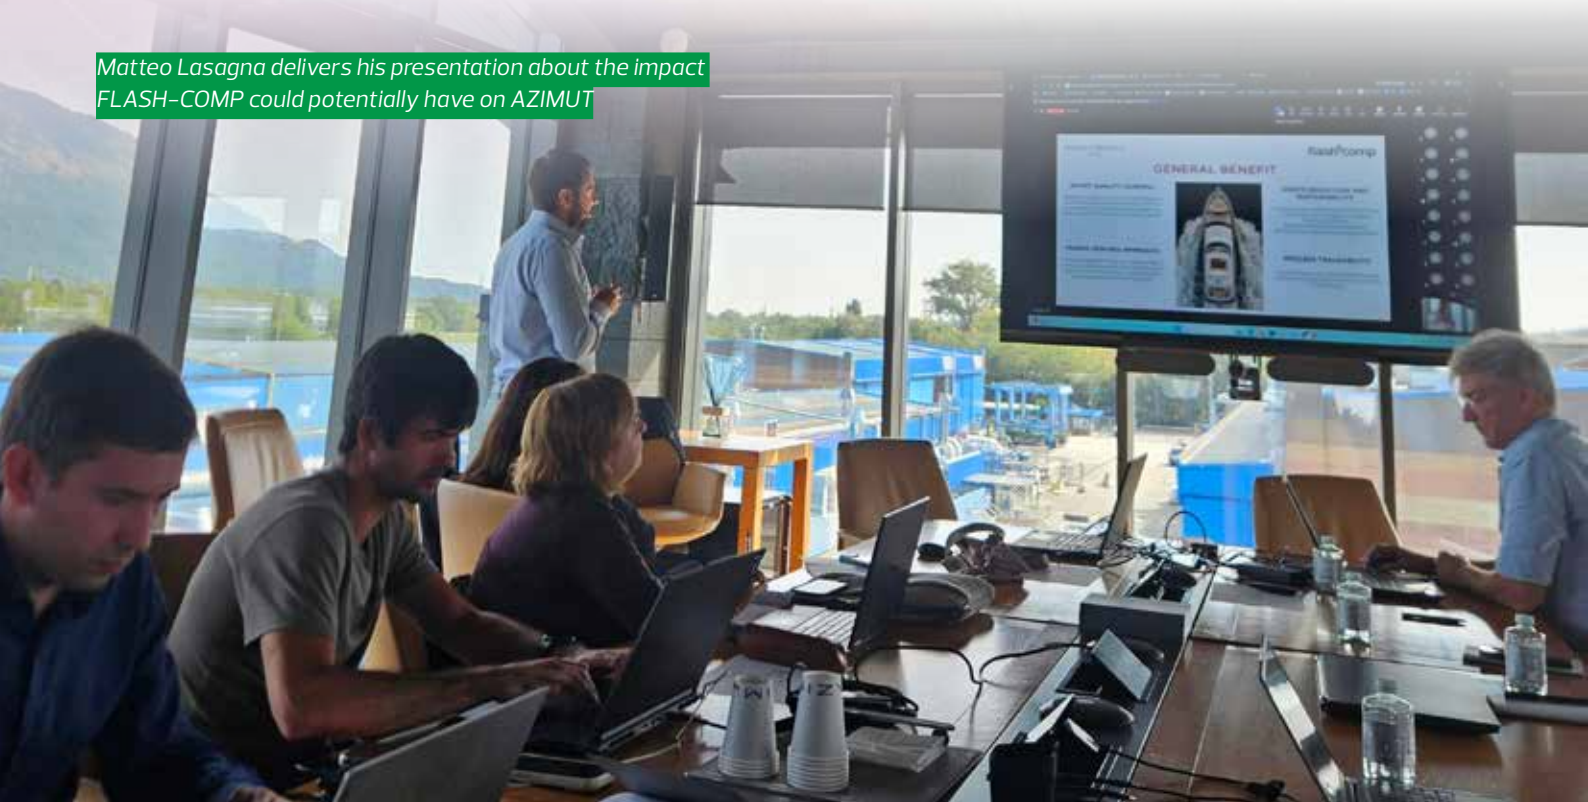
adoption remains unsatisfactory. Public programmes support pre-competitive development, but a gap opens when the technology needs private investment for productisation, certification, installation, sales and support. Small and medium-sized technology providers may not have the resources to cross that gap, and manufacturers may hesitate to invest before repeatability and return are proven.

Ighina argued that projects are often approved with expectations of high technology readiness that do not match their budgets, duration or instruments. Applied research should remain ambitious, but it should be frank about what can be achieved. A four-year project may deliver a relevant environment demonstration and a credible roadmap. It rarely delivers a fully industrialised solution across multiple lines and sites. Funding structures should recognise follow-on stages and allow progress to be supported step by step against evidence.

He also challenged how projects are formed. Research organisations often develop an initial concept, then invite technology providers and end users to supply suitable use cases. This can produce good research, but the end user's operational problem may not drive the project from the start. Manufacturers themselves need better feasibility studies before joining a proposal and a clearer internal roadmap for what happens after the pilot. Academic partners need closer exposure to industrial constraints, while programme designers need more realistic expectations of industrialisation.

Pantelelis agreed that the problem has persisted across decades and across funding systems. European projects gave his company access to major manufacturers, technical resources and visible demonstrations that would otherwise have been difficult to secure. Those opportunities helped build commercially useful capability. His experience was closer to an exception than the general rule, and further support is still missing at the point where a promising result must become a dependable industrial product.

Craamer added a replication warning. Pilots usually run on the newest line with the best connectivity and the staff most enthusiastic about innovation. A successful result there does not prove easy transfer to older machinery, other factories or less favourable working conditions. Manufacturers also see



Matteo Lasagna delivers his presentation about the impact FLASH-COMP could potentially have on AZIMUT


Matteo Lasagna


Nikos Pantelelis


Jose Pedro Ferreira


Marta Pinzone


Peter Craamer


Gorka Alvarez

their processes as unique, even when a supplier has a strong reference in the same sector. Bringing a new system into a brownfield plant is therefore a partnership measured in years, not a software installation. A mature commercial software tool can already cause major organisational change. A new platform with artificial intelligence, prediction and operator recommendations creates a much larger transition.

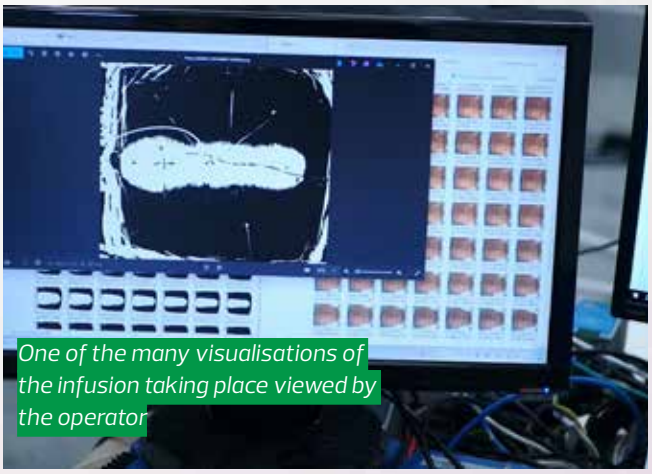
The panel did not regard a project as wasted when the complete integrated system does not go directly to market. Intermediate components, knowledge and methods can enter partners' products, services and later projects. The lesson is to assess these outcomes honestly and give each a suitable path rather than presenting every result as a near-market product.

Priorities for manufacturers and research programmes

The closing advice was practical. Manufacturers should begin with a strategic analysis of the problem and the defects that cause the greatest operational or commercial harm. They should collect clean, contextualised data, involve operators, assess digital maturity and build skills before purchasing complex automation. Early feasibility work should identify the component, production conditions, baseline, integration requirements and decision that the technology must improve.

Technology providers should minimise operator burden, make outputs explainable and adapt solutions to the chosen process and sector. Modular systems offer a better route than one fixed architecture because customers have different equipment, data environments and needs. Deployment should progress through descriptive, diagnostic, predictive and prescriptive functions only as evidence allows.

Research programmes should distinguish demonstration, productisation and industrial replication. They should support follow-on validation, include realistic readiness targets and encourage end user problems to shape projects earlier. Shared data spaces also need credible governance. Industrial data are commercially sensitive, and manufacturers will not contribute them merely because anonymisation is promised. Ownership, permitted reuse, confidentiality, cybersecurity and access after project closure must be explicit.


One of the many visualisations of the infusion taking place viewed by the operator



The welcome sign greeting the event delegates at AZIMUT HQ

Chapter 5

Exploitation and engagement after the project end

The exploitation deliverable translates these lessons into a portfolio strategy. It identifies five Key Exploitable Results: the Combined Decision Support System, the FLASH-COMP Data Platform, INEGI's Simulation Tool, hyperspectral cure monitoring and the fibre optic monitoring system. Five additional Exploitable Results cover the Defect Estimation Tool, Data Augmentation Dynamics, the INEGI testbed, the industrial demonstrators and the Manufacturing Sustainability and Zero Defect Manufacturing Maturity Methodology. Training modules, the Standardisation Roadmap and scientific publications form three Usable Results whose value lies mainly in continued use, capability and knowledge transfer.

The Combined Decision Support System is the clearest integrated opportunity. MSI's generic data architecture can receive, process, store and share industrial information. INEGI contributes composite process knowledge and simulation, while LORTEK contributes analysis and decision support. Monitoring technologies supply the process evidence and the human-machine interface turns the result into operator information. The digital chain was demonstrated with Azimut data, including real-time communication, standardised alerts, historical replay and separate live and training modes.

Further joint development will be based on shared ownership between MSI, INEGI and LORTEK and is confirmed in principle. Precise shares, exploitation leadership, delivery responsibilities, support, liabilities, revenue allocation and data rights still require agreement. INEGI has been proposed as a possible lead exploiter because of its composite manufacturing expertise, with MSI and LORTEK as technology providers. That arrangement should proceed only when the partners identify a viable customer, funded opportunity or strategic commitment.

The portfolio is likely to reach users through services, integration and licensing before standard product sales. Paid pilots can test the Combined DSS in a defined production case. The data platform can support MSI led integration projects and cross-sector replication. Simulation may be delivered as an engineering service or software component. Fibre optic monitoring may suit licensing to a specialist integrator. The maturity and sustainability method can support assessment and advisory work. The testbed and demonstrators can generate evidence, host trials and support training. Across the portfolio, the next development phase is industrial validation with common measures. Trials should record defect avoidance, scrap, cycle time, labour, energy, integration effort, usability and total cost. Benefits must be tied to a clear baseline and should not be promoted as confirmed savings until the evidence supports them. The hyperspectral route, for example, needs lower-cost hardware and broader validation.



The blue vacuum bag on the piece being infused and monitored by FLASH-COMP technology

AI and defect estimation need sufficient labelled data. The wider monitoring package needs repeatable installation and sector-specific configuration.

The targeted engagement strategy provides a route from communication to exploitation. Insight Media is matching each result with relevant manufacturers, integrators, researchers, training providers, clusters, standards stakeholders and validation partners. Responses, barriers and sentiment are recorded by result. A contact becomes a qualified opportunity only when a credible use case, trial, collaboration or reuse pathway exists. The records will then be handed to the relevant result owner.

This work has already produced identifiable continued-use outcomes, including use of training materials by an Italian educational institution, application of the six dimension maturity process and conversations with companies interested in Zero Defect Manufacturing.

A phased route to continued use

The evidence from the event and the exploitation plan points to three connected phases, which could be used as template for other research efforts considering their post-project further use strategies.

Immediate definition and evidence

Result owners should confirm ownership, access and the minimum offer for each priority result. For joint solutions, partners should agree leadership, contracting, delivery and support. A small number of industrial cases should be selected with clear baselines and decision criteria. Targeted contacts from the engagement work should be converted into technical discussions, pilot specifications and validation partnerships.

Further validation and productisation

The partners should test repeatability on additional geometries, resins, lines and working conditions. They should quantify the cost of installation and operation, reduce hardware cost where necessary, improve integration with existing systems and record operator experience. Training and the maturity assessment should accompany the technical deployment. Paid pilots, contract research and follow-on public funding can support this phase when direct investment is premature.

Replication and sustained use

Solutions that demonstrate value should be packaged as modular services, licences or integrator supported deployments. Reference customers and demonstrator evidence can then support expansion within marine and aerospace before less evidenced sectors are pursued. Training, publications, standardisation knowledge, testbeds and methodologies should remain accessible with named



A FLASH-COMP visualisation of an infused part shown during the LRI process

custodians, update arrangements and measures of use. Results that do not support a direct market offer can still create substantial value through research, teaching and future collaboration.

Conclusion

The FLASH-COMP event showed how sensing, simulation, data and decision support can move quality control earlier in liquid resin infusion, then tested that approach against industrial reality. The Azimut validation connected preparation, infusion and cure. The marine and aerospace talks showed its adaptability, while the readiness, sustainability and skills work showed that adoption depends on the organisation as well as the equipment.

The roundtable's central challenge was that industrialisation requires time, investment and responsibility beyond the demonstration. Clean data, realistic feasibility studies, operator involvement, brownfield integration, repeatable evidence and a credible financial case determine whether technology leaves the pilot environment.

FLASH-COMP enters its post-project phase with a credible portfolio and clear decisions. The most promising immediate route is modular and customer-specific, led by paid pilots, services, licensing and further validation. Joint solutions need formal governance, and each result needs evidence proportionate to its claims.

The project's lasting contribution may therefore be measured in several ways: technologies advanced into industrial service, methods and training adopted by others, data and demonstrators reused in new work, and partnerships that carry the results into further validation. The event showed that these outcomes are possible. It also defined the conditions required to achieve them, giving FLASH-COMP a practical route from an industrially validated research project towards wider manufacturing value.



Komola Munda collects clean water from the rainwater harvesting system in Kolbari, Bangladesh. June 2023.

WaterAid/Fabeha Monir

Together, we make a lasting difference.

We are an international not-for-profit, determined to make clean water, decent toilets and good hygiene normal for everyone, everywhere.

Get involved at **wateraid.org**

WaterAid

Registered charity numbers 288701 (England and Wales) and SC039479 (Scotland)