The development of future large-scale safety-critical systems, also known as cyber-physical systems (CPS), involves many technology and influencer communities. Novel approaches and tools will be required to tackle the multi-dimensional challenges between these communities to benefit CPS, especially their adaptability to new technologies.

Bridging the stakeholder communities that produce cyber-physical systems

By CHARLES ROBINSON ET AL.[1]

There are many communities involved in the creation of cyber-physical systems (CPS), which are used in domains including transport, health, manufacturing and, in the longer term, will be in the home, where miniaturization will play a role.

In this article we explain that engineering for future CPS needs a centre of gravity in order to draw these communities together. This will provide common goals around which technical advances can be aligned. Overviews of the communities involved are provided, with examples of their relevance in the creation of CPS and to some common challenges.

Advancements of aggregating technologies are multi-dimensional challenges, representing many influencing dependencies from all communities, especially at higher levels where the whole system product is drawn together. This means that, to make good progress, Europe will require new forms of coordination in order to orchestrate research and to capitalize on lessons learned related to the cumulative advances between the communities.

Key insights

- Large-scale safety-critical systems, also called CPS, are physically interactive (high certification obligation) and increasingly collaborative (task sharing). They involve many contributor and influencer communities in their creation, who each tend to make advances in isolation. Creating the technical bridges between these communities to channel technology development is essential for these future systems.
- The scope is wide and communities need a technical interface around which to align. Discussions suggest this centre of gravity to be real-time safe and secure automation.
- A new form of research coordination is necessary to direct cumulative developments from the stakeholder communities. CPS projects with cross-community challenges and which involve most of the stakeholders are needed for this.
- The development of CPS requires a holistic approach, guided by target products, that brings together a wide range of disciplines. These should include not only functional, system and enabling technologies, but also the fields of psychology, sociology and ethnography, among others.
- Aggregating technologies have different industrial uptake lifecycles to component technologies. Research programmes treat them the same and both technology types suffer. A dedicated team would be very beneficial for investigating and implementing specific technical supports (for projects, programmes and industrial policy).
- System thinking/interdisciplinarity will play a critical role in reversing the adverse effects of technology on the environment and requires more support from national funding. Technologies for future large-scale safety-critical systems is one key area that this applies to, especially to tackle their high environmental impact.

---

[1] For the list of contributors, see the acknowledgements at the end of this article

Key recommendations

- Draw on best practices for systems thinking/interdisciplinarity to develop methodologies and supportive tooling that will act as the technical bridges between communities contributing to large-scale safety-critical systems.
- Real-time, safe and secure automation should be used as the centre of gravity for communities contributing to large-scale safety-critical systems.
- While advancing towards this centre of gravity, we should not neglect the other common interests (common challenges) shared between the communities that will reinforce the bridges. Approaches from other interdisciplinary fields should be used as inspiration.
- Research orchestration needs to be developed for coordinating cross-community research to future large-scale safety-critical systems. This also calls for projects tackling challenges that are common across the contributing communities.
- A research instrument, transversal to European research programmes with stable financing, is required for advancing aggregating technologies in particular and the technology uptake by large-scale safety-critical systems in general.

Introduction and new cross-community development approaches

In order to manage large complex problems, people break them down into parts. It is for this reason that, from the technology point of view, there are many contributing and influencing communities involved in the creation of future safety-critical products. Of course, the parts subsequently need to be assembled together in order to address the initial complex problem. For the same reasons, the various technological contributions for future large-scale safety-critical systems, also known as CPS, require layered aggregation in order to achieve these physically interactive and collaborating systems.

This means that there are significant, multi-dimensional influences across CPS communities, which contribute to our ability to transfer technology to industry. It also poses challenges for assuring CPS, which should be based on sound methods of justifying that a CPS is fit for purpose and that all risks of using it are adequately addressed, notwithstanding the complexity and the heterogeneity of the CPS components and of the communities of stakeholders involved.

For the purposes of this article, we take CPS in the context of an application; that is to say, the term could be replaced directly with an example CPS application such as railway transport or satellite constellations. In this framing, CPS therefore represent physically interactive and collaborating systems that are present in many domains including transport, health and manufacturing (For an in-depth definition of CPS, see, for instance, the HiPEAC Vision article “Understanding cyber-physical systems among many communities: Large-scale safety-critical systems”).

Communities involved in CPS, discussed in the subsequent section, range from providers of a) functional properties including sensing, physical action, communication, energy provision, processing and coordinated collaboration to b) system-level engineering including properties like safety and performance specifications, managing customer requirements, architecting, system validation, mechanical engineering and control engineering. There are technology support communities providing c) enabling technology domains like the Internet of Things (IoT), Systems of Systems, Big Data, Artificial Intelligence (AI) / Machine Learning (ML) and High-Performance Computing (HPC). Finally, there are the influencing communities from d) the production environment, with enterprise processes and product line, and e) the market, such as regulation and current and future needs of society.

These communities have tended to transfer technology as a one-to-one mapping with products. However, to respond the challenges of future CPS and to enhance technology transfer, they will need to take relations with the other contributing communities increasingly into account. While the challenges and importance of advancing aggregation techniques are discussed later, there also needs to be a common focal point from which one community can interact with any of the other communities. This point should provide a common interest based on the physical challenges of these systems. *Discussions have proposed this centre of gravity to be real-time, safe and secure automation of CPS development and operation.*

Research on CPS should seek to enhance the interrelations and automation of these three dependability properties, i.e. real-time, safety and security. They are goals that must be achieved at a global level when all the technologies are combined. As an example, each piece of hardware has an impact on the energy consumption of the whole system. Similarly, individual software and hardware components can jeopardize safety if they fail naturally or due to a security breach. These goals can also be variable and related to environmental conditions, such as a train reducing its speed (performance) in response to heavy showers (to maintain safety).

Hence for technologies to be accepted in these systems, they must guarantee these dependability properties, i.e. they must comply with the safety and security constraints of a product and not compromise real-time responses. This means that the easier it is to couple your technology with these system constraints (through automation), the easier it becomes to adjust it to the system (or adjust the system for new technologies).

It is usually the case that, in order to add new technologies to a CPS, the whole system requires re-certification. This can be prohibitively expensive without sufficient automated information about the impact of the new technologies on these dependabil-

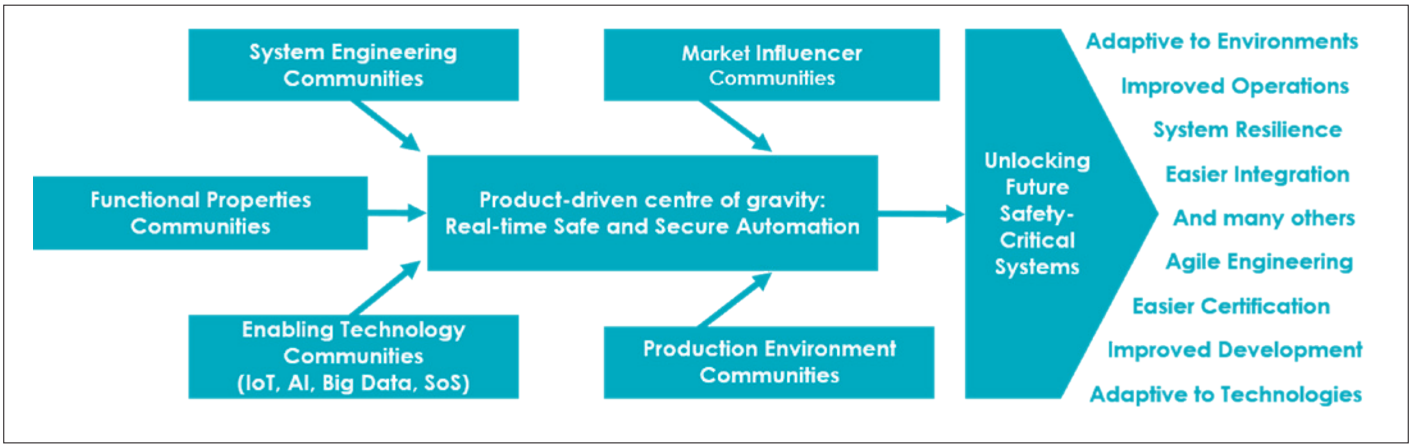

Figure 1: The stakeholder communities for creating CPS.

ity properties – and particularly the interrelations of those properties. Take systems certified, for instance, against an extreme earthquake occurring every 1000 years (for safety), such as a nuclear plant: in this case, the safety experts currently would prefer no new technologies or patches for security to be added to these systems due to the certification costs.

As a result, historically, interrelations between system properties have been limited to minimize complexity, but the current need for adaptability (to new technologies, to environmental or internal changes) requires this design mindset to be readdressed. So in summary, a centre of gravity, as shown in Figure 1, will provide a useful point to channel us towards more impactful research advances for these future large-scale safety-critical systems.

While the management of trade-offs between the system properties of performance, safety and security is an established skill in system development, it still remains very much a manual and qualitative process and one that is based on prior experience. It remains to this day very much a bottleneck and is holding back the communities contributing to CPS development from ensuring that advances in areas such as trust in artificial intelligence (AI) are applied to CPS.

System-level engineering for CPS is therefore in need of transformative automation. Fortunately, automation between system-level dependability properties can rely on a number of decades of research in techniques [1], some of which have already been applied in industry but are generally in need of new approaches for technology transfer. Such approaches are included in the coordination suggestions for research orchestration described later in this article.

Of course, current pressures for industry to find advanced solutions for managing system property trade-offs are also driving the search for automated coupling. As examples of some initiatives, the UK Research Institute in Trustworthy Interconnected Cyber Physical Systems (RITICS) [2] involves dozens of UK universities and industrial collaborators. Topics include safety, security and autonomous systems. Relating to autonomous vehicles, the Intel Research Collaborative Institute of Safety of Autonomous Cars (ICRI-SAVe [3]), deserves a mention as a vibrant community.

Many industries are actively looking for solutions to manage the performance, safety and security of their products, including large enterprise like Siemens, Thales and AVL, who have been forming combined safety-security teams. The challenge also affects small and medium-sized enterprises (SMEs) in their products and services. This recent momentum has visibility, for example, in the Ada and IEEE conferences, in the IET code of practice on cybersecurity and safety [4], in recent large research collaborations including MERgE [5], SeSaMo [6] and AQUAS [7], and in co-engineering discussions.

Overview of stakeholder communities for creating and advancing CPS

We now provide overviews of the five communities, indicated in the previous figure, which are involved in creating CPS. We give descriptions and examples of their relevance to CPS as well as their relation to cross-community challenges for future development. These include embedded computing as a CPS backbone, system decentralization and decomposability, and physical collaborations with people.

Functional-property communities

CPS functional properties have to address aspects that cover sensing, actuation, communication, energy provision, processing and coordinated collaboration. Such properties are key characteristics of these systems, with actors in specific communities researching and developing the different components.

The relevance of functional properties becomes more evident when considering novel and innovative advanced applications that are being progressively adopted in a number of large-scale, safety-critical domains, such as industrial automation, transportation, smart cities, critical infrastructures, space, etc. Some examples can be found in H2020 projects such as CPSwarm [8] and other CPS cluster initiatives.

Industry-driven needs and the well-established nature of general research communities in the CPS domain mean that it is feasible to envision projects that might prototype concepts such as swarms of unmanned aerial vehicles and rovers supporting safety and security operations; swarms of automated ground robots that collaboratively support humans in logistic operations within a smart warehouse; or enhanced and dynamic platooning applications for autonomous freight vehicles. Currently, the development of such applications cannot leverage a simple plug-and-play integration of the various technologies entailed, given the complexity of managing teams of systems and humans in evolving and dynamic scenarios with emergent properties.

Therefore, in order to properly combine and integrate the different technology building blocks required, the various 'functional properties communities' have to be properly engaged. Experts from the functional property communities will need to work with other actors with collaborative systems competence. Moreover, while the increased adoption of CPS has resulted in the maturation of solutions for CPS development, a single consistent science for future CPS has not yet been consolidated. Few functional properties community members have already started working alongside other communities on a connective framework e.g. using modelling, design/development tools and methodologies, deployment solutions, monitoring and controlling solutions for large-scale challenges. In this context, model-centric approaches have clear relevance for facilitating collaboration between experts from different sectors and thus enabling the definition, composition, verification and simulation of collaborative, autonomous CPS.

For these reasons, it is important for future CPS to be considered not only from the technology perspective but also as an **application domain** where the technology of the functional properties' community plays a role for aggregation of CPS-related research. To promote this, closer and wider collaboration is needed within the community, along with new research initiatives. Understanding the nature of this aggregation from the bottom up and top down is important for driving the communities towards much-needed technology advances. The resulting collaboration plays a very important role in finding solutions to the bottlenecks that currently prevent CPS from having greater impact on society; such solutions would also promote market uptake, open up new markets and optimize the use of resources in the various industry sectors.

These communities have many cross-cutting challenges for future CPS. Embedded computing will evolve significantly and plays an essential enabling role for functional properties. For instance, the need to use specific sensors on a CPS and to timely process the relevant raw data onboard will need increased computational power. However, energy limitations introduce other constraints; only a holistic vision of CPS can help driving research initiatives. Moreover, the envisioned combination of 5G, beyond 5G and Smart Networks and Services/6G technologies with distributed and high-performance computing will pave the way towards a deep integration of future CPS in the computing continuum. In relation to decentralization and decomposability, with distributed intelligence and emergent properties, an example research context would aim to solve/work on delays in physical, computing and actuation timing. This requires model design and simulation approaches to capture the full heterogeneity of the system and its contributing communities. Physical interaction with people requires a system to have high fidelity knowledge of its environment and its physical dynamics. This requires the technologies of the functional properties community, which in turn need integration with the safety and security measures set by the system-engineering community. It is therefore clear that the best way to advance future CPS is to further support integration and aggregation approaches for community collaboration.

Systems-engineering communities

The development of CPS requires a holistic development approach that brings together a wide range of disciplines. This includes the typical systems engineering disciplines, such as requirements engineering, architectural design, implementation and quality assurance including system-wide responsiveness, safety and security. The disciplines of this community are important in terms of both the CPS in general and individual systems engineering sub-processes, such as mechanical engineering, control theory, electrical engineering and software engineering.

In almost all of our application-driven future scenarios, like in autonomous driving and Industry 4.0, CPS must be able to fulfil their purpose to a large extent without intervention of human users [9]. According to the Society of Automotive Engineers (SAE) taxonomy for autonomous driving, we refer to such systems as highly automated or fully automated CPS [9]. Already today and even more so in the future, systems engineering is one of the core competence fields for building such highly automated or fully automated CPS.

In the case of highly automated CPS, it is necessary to have a more comprehensive understanding of the term 'functional safety'. In contrast to the understanding of the term by the ISO 26262 standard, which essentially considers the malfunction of system components, highly automated CPS require an analysis of the interaction of a) the functionality of the CPS under consideration with b) its context (e.g. other CPS in collaboration). This analysis serves to detect possible safety threats resulting from the interaction between system functions and contextual conditions, such as the interaction between the autonomous driving function of a vehicle and the failure of the signalling system at an automated road intersection. This new understanding of functional analysis, which goes far beyond the requirements of ISO 26262, is the subject of the SOTIF standard [10].

These threats to safety must be identified during the development process and mitigated, e.g. by specifying suitable requirements or safety devices (safety monitors) which bring the CPS to a safe state should CPS fail to behave according to requirements/expectations. Since CPS often monitor and control technical or physical processes, control theory is a discipline of great importance in the development of such systems. In this context, the concepts of monitoring and controlling technical/physical processes are reflected in various artefacts of systems engineering. For instance, the requirements originated from the way the processes should be controlled, as well as from decisions made about the design of the necessary sensors and actuators or even about the design of the algorithm for the computational processes of the feedback system.

In order to be able to develop such complex technical systems consisting of software and hardware, seamless systems engineering processes are required, establishing techniques, methods and tools for challenges such as the following examples. Since CPS in many fields of application work together in dynamically formed networks at runtime to pursue higher-level goals, possible collaboration structures must be identified and analysed in requirements engineering. For example, in the development of autonomous vehicles, the collaboration structures in which these vehicles must operate should be taken into account. Examples of such structures might be vehicle convoys to optimize the flow of traffic or at automated intersections to ensure safe crossing of the intersection, even with high traffic volumes and in complex traffic situations. In collaborative CPS, the issue of coordinated decentralized monitoring and control of technical/physical processes is added; an example of this is the coordinated acceleration or deceleration of the various vehicles within a convoy of vehicles.

In the case of highly automated systems, the involvement of the human user is required in (a few) defined situations to ensure that the system is able to fulfil its purpose of ensuring safe operation. The integration of the human user must be effective, i.e. the user interface of these systems must be designed in such a way that the human user is able to perform the necessary tasks according to the intention, as free from errors as possible and within the existing time restrictions. One might think here of the example of autonomous road traffic, where highly automated systems require the driver to take control of the vehicle when a critical driving situation occurs.

Enabling-technology communities

Internet of Things (IoT)

The Internet of Things community developed around the goal of providing a means for all devices to be globally connected via the internet. The name 'Internet of Things' was used in 1999 by Kevin Ashton during a presentation to his higher management at Procter & Gamble. He described IoT as a technology that connected several devices with the help of RFID tags (radio frequency identification) for supply-chain management [11]. In 2008 the first international conference on IoT took place in Switzerland, discussing RFID, short-range wireless communications, and sensor networks; today, these topics continue to represent the major technological research domain for advancing the IoT, gathering information about the real world that can then be made useful in some way [12].

Since 2010 it has been normal for many different devices to be in our homes to be connected to the internet. Connected devices are used extensively in the consumer domain. In 2015, to support advancement of IoT for industry, the European Commission created the Alliance for Internet of Things Innovation (AIOTI). Applying IoT to the industrial environment has been termed industrial IoT, or IIoT, and has the goal of optimizing production value while considering the many additional challenges related to safety, security and performance. IIoT technologies support interconnectivity with the internet in the context of these challenges, enabling not only networked smart objects and information technologies but also "*optional cloud or edge computing platforms, which enable real-time, intelligent, and autonomous access, collection, analysis, communications, and exchange of process, product and/or service information, within the industrial environment*" [13]. IoT technologies , in particular those for the IIoT, will be standard constituent elements of future safety-critical frameworks.

Enabling the infrastructure to support distributed intelligence and information exchange is at the core of IoT, so supporting cross-community work on CPS decentralization, decomposability and human interaction is important. These are already areas receiving some focus from the IoT community [14, 15], as indeed is the case for bringing communities around an embedded computing backbone, with work considering edge-cloud computing [16] exchanges. As an enabling technology, IoT responds to support other domains which means its focuses change based on the latest domain challenges, corroborated in recent IoT roadmapping activities that its landscape is changeable in nature [17].

Artificial Intelligence (AI)

Autonomy will bring incredible new benefits to CPS, but there are major challenges that must be overcome. The intelligence that can be applied is limited by current approaches to certification, legal frameworks and (lack of) trust for such systems. There is also a fundamental mismatch between the approach to functional safety software (top-down, correct by design) and current approaches to deep-learning programs (bottom-up, data driven), as addressed by projects such as SAFEXPLAIN [18]. These need to be addressed while maintaining and increasing the safety of such systems (which calls for improved traceability of the influences between the contributing communities to CPS). Safety of such systems is a serious challenge as the levels of reliability achievable by ML/AI are simply inadequate for high integrity systems (safety integrity levels 3 and 4) by and are an active area of research attracting investment from public and private sector. Cyber security of such systems is even greater challenge as vulnerabilities of ML/AI are significant and widely spread.

Reducing or mitigating these limiting factors will be an enabler for many advanced AI technologies related to decision making, learning etc, for the operation of the systems. In parallel, the other communities can provide more robust technologies for systems that are evolving as a result of AI. Of course, there are identified routes for AI to become "more trustable"; these include explainability of actions in human language, and the application of AI to non-safety-related aspects of CPS like decision support for system design.

A significant characteristic of CPS will be coordinated collaboration. This relates to the way components of a CPS coordinate with each other or with people for outcomes only achievable through such cooperation. AI can bring strong support here such as through the field of decentralized intelligence called multi-agent systems (MAS) [19]. Regarding design, the needs of CPS include the explicit representation of the environment and the need to represent abstraction layers, from the physical layer to the components and system, as CPS are closely coupled to the hardware elements of the system. Finally, it may also be necessary to represent the non-functional requirements, such as safety or resilience. Some MAS design tools, such as Tropos [20], if correctly used, may help to meet these requirements.

In terms of decentralized intelligence for CPS, there are many challenges to that need to be addressed, in particular methods for executing coordination. The whole system needs to be able to react in real time, which is not the case for most decentralized AI coordination protocols, which rely on negotiation, usually with no defined deadline for decisions [20]. As another example, finding ways to work with the functional property community on communication middleware for intelligent collaboration is likely another issue to needing to be tackled.

High-Performance Computing (HPC)

High-performance computing (HPC) consists of the aggregation of powerful computing resources for solving problems that require large computing power [21]. Recently, HPC technologies were only required in the context of traditional massively parallel “number crunching” applications like weather prediction, computational chemistry, or computational fluid dynamics. However, the latest developments in low-power computing technologies [22] – required in the HPC industry to scale performance levels further – has facilitated the adoption of HPC technologies in a wide range of CPS applications.

Existing HPC platforms offer the computation capabilities needed by the most demanding CPS applications within an affordable power budget in domains such as automotive, space, avionics, robotics and factory automation. Centralized domain architectures that replace the traditional federated computing architectures – like those required by economically affordable autonomous driving systems – are only possible when HPC technologies are deployed. Single-chip high-performance embedded computing platforms reduce the traffic flow through CPS’ electronic networks and enable high-speed communication as required for processing vast amounts of information in real time. So this community will be important for consolidating the embedded computing backbone.

Furthermore, these technologies involve parallel processing, that is, splitting the tasks up into parts for several computers (or multiple cores) to process, thus reducing the time taken to complete tasks. This characteristic thus holds a direct relation with the CPS challenges of decomposability and decentralization – how tasks can be split up while ensuring safety and security for people, the system and its environment.

Unfortunately, the deployment of HPC in a CPS increases the complexity of the resulting system and may have a non-negligible impact on the verification and validation costs of relevant system properties (e.g. safety and security). Thus, an effective exploitation of HPC technologies in cyber-physical applications requires at least either the development of new methodologies to verify and validate such complex systems or the adaptation of key technologies to the specific context, as explored in the EU-funded PROXIMA [23] and MASTECS [24] projects, for example.

Big Data

Cyber-physical systems are being driven by the combination of embedded and internet technologies and a vision of “smart anything everywhere” [25]. The blend of this cyber, physical (and social) data can help us to understand incidents and changes in our adjacent environments better, monitor and control buildings and urban infrastructure, and provide better healthcare and care services for older people, among many other applications. To make effective use of the physical-cyber-social data, integration and processing of data from a variety of heterogeneous sources is necessary. A key objective for big data in CPS is to analyse very large, fast, and heterogeneous data streams, mostly from industrial rather than consumer environments. This can be achieved through machine learning, which is the most common technique used to extract information from the data.

The core Big Data applications in CPS are in varied fields, including energy utilization, city management, transportation systems and disaster management. For example, a smart transportation system would generate big data consisting of drivers’ behaviour, commuter information, vehicle locations, traffic-signal management, accident reporting, automatic fare calculations, and so on. Robot-aided surgical systems (i.e. human-in-the-loop CPS) comprise a teleoperation console operated by a surgeon, an embedded system hosting the control of the automated robot, and the physical robotic actuators and sensors. Big Data methods can be used here for the modelling of surgical skills, for the detection and classification of surgical motions for automation and environment, and for the integration of this knowledge into control and automation of surgical robots.

In the operation of complex systems (e.g. aircraft and industrial processes), fault-detection and fault-isolation schemes are designed to detect the onset of adverse events. Such systems use big data methods (such as machine-learning classifiers) to enhance the diagnostic accuracy of the online reasoner on board the aircraft. Moreover, big data can be utilized in command and control with cyber-physical infrastructures for emergency services and defence.

The value of the Big Data community as a contributor to CPS products can only grow in the future due to increasing interest in data as an important business asset. The combination of heterogeneous data from numerous sources will require new applications for integration, query and analysis, along with embedded computing, HPC, and data-reduction techniques. This remains an open research issue for CPS. The variety of types and sources of data will give rise to new kinds of data stores to sustain flexible data models.

Another important issue is that of remote storage of big data. Until now, cloud-based

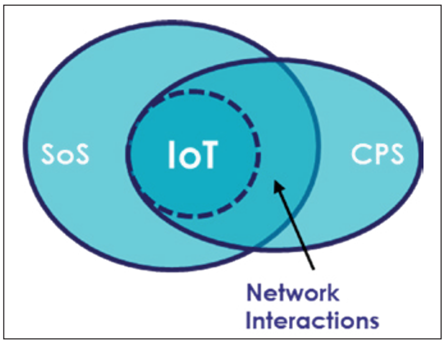

Figure 2: Technology relations of SoS, IoT, and CPS [35]

models have facilitated the storage and processing of big data sets, providing data accessibility and better IT power. However, this creates a centralized data store that does not scale in the CPS setting. To facilitate decentralized data storage and processing, a number of problems (e.g. replication, parallelism and requirements) arise. There is an urgent need for new approaches and techniques.

System of Systems

The "System of Systems" (SoS) concept has been around for at least fifty years, but in the last twenty it has been an area of major concern. Following the description of its characteristics by Maier [26]; it is defined in ISO15228 as: "*SoS…brings together a set of systems for a task that none of the systems can accomplish on its own. Each constituent system keeps its own management, goals, and resources while coordinating within the SoS and adapting to meet SoS goals*" [27]. As for CPS, SoS represents a type of application as well as a technology domain.

Broadly, one can consider SoS applications as independent systems that interoperate (work together) to achieve a purpose, with a significant amount of ubiquitous networking. In the case where they have extensive software control between safety-critical systems, the application itself is both a SoS and a CPS because they share common characteristics. Figure 2 describes the relationship between SoS, CPS, and the Internet of Things. Where infrastructure interactions are supported by internet protocol, then the CPS is also described as IoT, which is necessarily always a SoS. There are also interesting SoS-CPS applications that interact through means other than the internet protocol (e.g. mechanical or electromagnetic interactions) and the engineer may need to guard against such interactions for safety or performance reasons.

However, from the technology perspective, CPS application research considers how all technology communities are integrated to create a system and its interactions, with the SoS technology community contributing to the coordinated collaboration aspect. This is a key property for future CPS, meaning that SoS research is indispensable for creating future CPS. In relation to embedded computing, the importance of localized processing, while maintaining a connection to centralized processing capacity, is recognized as a priority in areas such as edge computing, which uses SoS technology. This also links directly with the challenge of decentralization or decomposability where systems work together. A smart city is an example of human interaction and SoS, for example; it manages busy traffic at city junctions to minimize delays for drivers and pedestrians.

In 2012, INCOSE conducted a survey to identify "pain points" for SoS practitioners, i.e. the problems that kept systems engineers and managers awake at night [28]. The study indicated seven main areas of concern: SoS authorities; leadership; constituent systems; capabilities and requirements; autonomy, interdependencies and emergence; testing, validation and learning; and SoS principles. It is no coincidence that creating CPS includes these pain points, because they are concerned with networked, intelligent systems of high complexity. This suggests that the communities of SoS and CPS have areas of common interest suitable for collaboration.

Digital Twins and the Metaverse

Human-CPS interaction will also advance with the advent of digital-twin and "metaverse" technologies [29], in particular when CPS operate in close proximity or hand-in-hand with human operators. The metaverse will provide haptic feedback over robots that complement and advance human capabilities [30]. Human operators will receive visual guidance in their view of augmented reality, and will obtain the ability to project themselves into the CPS they control. They will sense, act and interact through the impersonated system with other humans and with the environment in which the CPS operates. They will receive extended cognition and operating capabilities over swarms and manage the complexity of CPS hierarchies with ample application areas. Human caretakers may intervene in case of emergency or when service robots hit the boundaries of autonomy.

Swarms will act in harsh environments on Earth, in space and on remote celestial bodies instead of exposing humans to the risks they have to take today. Examples include mining, nuclear-waste handling and reactor deconstruction, but also asteroid mining and exploration. Replacing the internet with a network of immersive virtual worlds, cyber-physical systems will allow the metaverse to bridge into reality, with all the benefits, but also all privacy, safety and security risks this entails.

Digital twinning is one of the enabling technologies for exercising such advanced control from the digital realm over the real, physical world. Digital twins are virtual models of reality that are continually updated about the actual state of their physical counterparts and which can enable decision-making that, in turn, leads to changes in the real world. The long-term goal of digital twins is to be able to capture the intentions and objectives of the physical twin, but also to improve overall performance through digital simulation, testing and monitoring how the real-world physical system will act in its environment. While the aim is to advance into a better future, this can threaten safety and security when not handled with utmost care. Thus, it will be inevitable for the metaverse and digital-twin communities to join forces with the CPS community to achieve real-time safe, secure, and cyberattack-resilient automation from the moment metaverse-enlightened CPS are designed and throughout their lifetime.

Production-environment influencer communities

Members of the production-environment communities are responsible for the industrial product process and lifecycle. This includes enterprise policy and processes, decisions about technology usage and the evolving physical plant [31]. They drive the large-scale production of goods using equipment in the form of modular automated product lines. Such equipment typically combines mechanical, electrical, and software components; it also requires substantial initial investment and maintenance costs. Throughout its long lifecycle (15-30 years) [31], the equipment operator and component suppliers cooperate to repair and repurpose/upgrade parts at a minimal cost. This imposes several constraints on component models and their versions, which in turn constrains policy and process management.

In addition, the arrival of digitalization and the CPS revolution brings the "servitization in manufacturing" opportunity, a paradigm shift where manufacturers shift to offer product-related services, beyond just selling a tangible asset (For further discussion of this concept and examples, see the article "Everything as a service" in this HiPEAC Vision). In the above example of automated product lines, component providers could offer monitoring, online maintenance, repair, and overhaul services [32] among other value-added services. Service contracts generate more steady revenue compared to the cyclical product business, but, in general, organizations in manufacturing struggle to drive servitization [32], because the introduction of the new services incurs higher costs without proportional returns.

The adoption of digitalization tools and solutions and the development of innovative services leveraging the full potential of CPS require incentives and coordinated efforts among different partners. Research projects, partnerships in which early movers and less-digital companies cooperate to embrace servitization and adopt CPS tools, provide a nurturing environment, where decision-makers find that the "test-before-invest" concept is an incentive that helps lower barriers and can evaluate potential benefits. For example, in the H2020 HUBCAP [33] project, less digitally focused SMEs were able to pair up with model-based design providers to adopt digital innovation and enhance their solutions using model-based design technology.

Among the success stories, there is the example of the partnership between Mototok International GmbH, a provider of innovative aircraft tug solutions, and Evitado Technologies GmbH, a provider of LiDAR-based algorithms adding advances from the self-driving car industry to an already innovative CPS product. Other examples show how advances were made in training for industry 4.0, the development of innovative organ-preservation devices in the medical domain, smart textiles, and precision agriculture.

The prime innovative aspect of HUBCAP is a web-based collaboration platform that facilitates stakeholders' access to computing resources and advanced CPS design and engineering solutions, by providing a cloud-based sandbox solution (Figure 3). The sandbox provides pre-installed models and tools, allowing companies to experiment with new tools and assets in a ready-to-use virtual machine available via a regular web browser, with emphasis on performance and interaction between partners. This will be taken forward and combined with DevOps capabilities, also in a digital twin setting, at http://gitworks.io/.

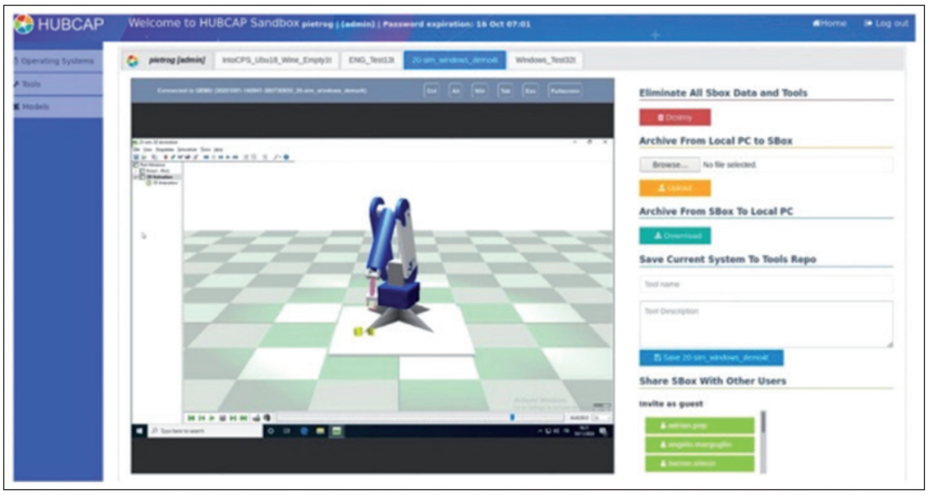

Figure 3: Snapshot from the Sandbox showing SME asset

Production-environment community members are deeply involved with the cross-community challenges identified. There is a historical synergy with the development and advancement of embedded computing, which will continue in the future. This community is always demanding advancements in embedded computing, and advances in manufacturing also affect how we produce the embedded platforms of the future. Regarding decentralization and decomposability, there are several lessons learned and case studies in which cooperation and adaptation to local and greener processes promote research, discussion, and changes to manufacturing. Finally, this community has a particular interest in the challenge of physical collaboration with people. This interest is from both an internal perspective, covering topics such as human-machine interaction and collaborative robots, and an external perspective, where the potential for improvement from product usage data needs to be fully explored.

Market-influencer communities (society needs, regulation, standards, policy)

CPS are believed to have an enormous impact on many aspects of socio-economic life. Therefore, a number of stakeholders grouped here under the generic name of 'market influencers' will have a stake in shaping the future of CPS and of the contributing communities.

Societal needs may be described basically by means of individuals or groups putting forward requirements and benefitting from CPS. The individual appears here as the consumer who is, in one way or another, making use of either a product incorporating CPS, or elements of larger CPS implementations, addressing communities of end users in terms of mobility, personal life (general wellbeing), healthcare, leisure, environment, etc. Other needs may be identified in the area of public services offered at local and national government level, including education, healthcare services, community services, and operation of public institutions.

As well as responding to societal needs, however, CPS also pose new challenges. Some specific fields include education and employment, as CPS induce the obsolescence of certain professions and create new ones. Therefore education, including training and retraining will be affected, as will the employability of the existing and future workforce, which will have implications for the labour market and social security.

Regulation – both hard and soft legislation - will have to be adapted in order to govern CPS so as to ensure their smooth integration into society. However, given the rapid cross-border spread of CPS technology, international agreements might be needed, too, particularly if we consider the globalised nature of today's value chains. Regulation will have to address the interplay between CPS actors (producers, consumers) as well the foreseen and unforeseen effects of the technology. Regulation is also supposed to be structured according to the societal needs that the technology is supposed to fulfil. A particular aspect of related regulation might address the human individual, chiefly in relation to human-machine interaction, which is anticipated to increase significantly in the coming years (intruding into both privacy and healthcare). The "must be implemented" regulation should be supplemented with recommendation-type measures of indicative nature.

Standards ensure interoperability and compatibility of products from different producers and allow the market presence of a large number of actors. Moreover, standards are important in order to set and describe safety levels and quality frameworks. To some extent, standards provide the technical base for legislation governing the area and also give room to innovation as usually standard specifications can be fulfilled in a variety of competing ways.

Policy aims to achieve certain results in a given field by reflecting society's needs or goals. Public policy in particular is directed towards supporting certain areas through frameworks of development in terms of tax incentives, grants or even regulation. Policy also includes public investment in facilities or processes of general interest. A further aspect for consideration is policies aiming to increase employment in a differential manner within the given population (i.e. in favour of disadvantaged groups), or to ensure development of regions lagging behind. Such policies also set out to address issues of general interest like climate change (that can only be done at international level) or the environment.

Beyond public policy, one should take into consideration policies of generically named "interest groups". Pressure groups such as non-governmental organizations (NGOs) consumer associations also have policies for their vision and procedures to supporting their realization, which can indirectly influence the market.

These "market influencer" stakeholders between them represent the conditions under which all the other communities operate for producing future CPS. The relevance of their involvement should be apparent, especially when considering the aggregative effects of contributing and cross-community technologies. Deficits in education in one community can have a knock-on effect on other communities. Training approaches and certification can be a deciding factor in the sustainability of mixed-community technologies. Policy can evolve approaches and perspectives that enhance behaviours supporting longer-term governance or culture, providing resilience, value generation and trust in new technologies.

Research orchestration for cyber-physical systems

With respect to coordinating CPS research *as an application domain*, additional approaches and orchestration should be introduced. This is because the application-domain perspective is based on the product side, with cumulative effects being considered through the aggregation of layered contributions from the stakeholder communities. Another issue is that disruptive discoveries, technologies or developments might influence the cycle of research. For example, if significant progress is made on quantum computing, or discoveries in material/biological science, that could make sensors more different.

Orchestration of research is particularly about knowledge management, longer development cycles, persistence and refinement of multi-disciplinary approaches for collaboration between communities. Take the example of constructing a building where a new team takes over every few months. Limited progress can be made without guidance at a higher level. This is similar for advancing CPS research. Persistence of acquired interaction techniques, between project collaborations, is significantly more difficult to maintain. For instance, usability and sensor experts have specific languages for their domains.

Therefore, approaches that support collaborations and which have been developed during collaborations should be taken, refined, and applied in subsequent collaborations of different groups. A dedicated CPS research instrument could advance this concept, in conjunction with future CPS support action projects. Projects themselves will also need to provide environments with favourable conditions for aggregative research considering the multi-

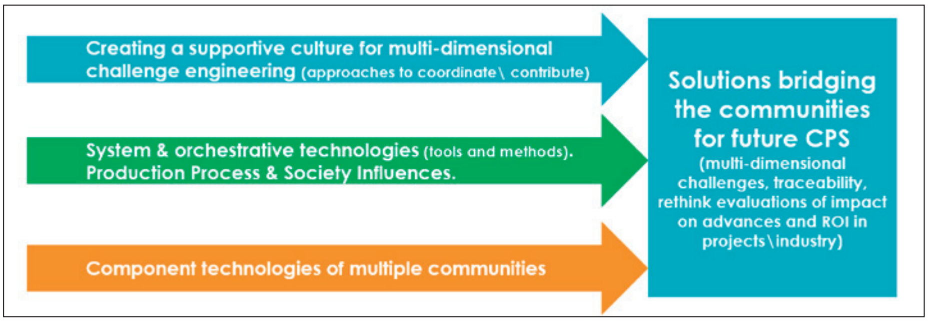

Figure 4: Stakeholder inputs to projects.

dimensional challenges, with conditions significantly different to those for developing component technologies.

Considerations for future CPS projects

For advancing CPS research as a technology domain, useful mechanisms already exist. For example, there have been projects following the standard approach, which gathers technology providers around one or more CPS-related use case. If awarded funding, the partners then work together for a few years to bring their technologies closer to market deployment (i.e. advancing "technology readiness levels" or TRL).

Cascade funding, where funded projects themselves fund smaller initiatives, has also shown itself to be a useful means for transferring component technologies for CPS, because the smaller initiatives are directly managed by companies looking for particular solutions.

However, *for the application domain side* of CPS research, new project approaches and higher support mechanisms also need to be introduced, enabling the multi-dimensional challenges previously discussed to be tackled. The characteristics that are believed to be essential in such projects are:

- Use cases: physically interactive and collaborative systems; of relevance to all communities, likely to be uniquely large industry or with integrated small-medium enterprises. Supplied also with the intention of advancing industry-side "industrial readiness levels" of production and product lifecycles for new technologies.
- CPS centre of gravity: all projects addressing the multi-dimensional challenges between communities should interface on work advancing real-time safe and secure automation, including interrelations, for CPS design and operation.
- Cross-community challenges: projects on application domain research should focus on grand challenges that need contributions from each community. Proposed call topics include:
  1) Embedded computing backbone
  2) Decentralization and decomposability
  3) Physical collaborations with people
- Developing the support environment: tools and approaches are required not only by industry, but also by researchers to support engagement of the different CPS stakeholders and perspectives. We propose that such projects include some dedicated work (a work package) that develops support for collaboration on the multi-dimensional challenges. New approaches established iteratively: orchestration approaches should be implemented in a manner that can be refined. We should avoid 'one-hit wonders' that seek to solve everything at once. A second iteration of such projects could also include smaller spin-offs and initial stage smart city investigatory projects.

Contributions of the communities to be these projects can be visualized, as shown in Figure 4, to involve the technology component providers, the influencers/aggregative technology providers and those developing the culture and support environment. This provides the means to advance the CPS aggregation techniques which are required to address the multi-dimensional CPS challenges.

These HiPEAC-proposed project characteristics (for application-domain research) relate directly to previous CPS community recommendations to the European Commission, including trustworthy and societal scale CPS, ethics data protection and liability, CPS engineering, interoperability, complexity, edge computing, humans-in-the-loop, co-engineering of system properties and enhancing uptake of CPS technologies.

Considerations for future advisory coordination and support actions for the CPS communities

The European Commission funds coordination and support actions (CSAs) to accompany, coordinate and stimulate innovation in particular technology fields and their communities. A particular challenge for a CPS CSA is that it is in fact a multi-community subject. This is because, as discussed, while CPS is *a technology domain* with specific complex challenges related to cyber and physical integration and cyber to physical plan realization, CPS is foremost an *application domain*. This is of consequence because CPS and other technologies can be much more difficult to apply to the final systems without also advancing the means for their combination.

To support application domain research projects, future CPS CSAs will likely support the transfer and synchronization of project environments, support the "big picture" metrics of aggregations in CPS and specific return-on-investment (ROI) valuation techniques to pre-empt industry needs. In particular, they will support a focal point for all the contributor technology and influencer communities.

Research instrument for technology orchestration: Supporting projects, funding programmes and industry

A team providing support across relevant research programmes, dedicated to supporting the channelling of different community contributions into safety-critical applications, in particular, would be very beneficial for supporting in particular technology orchestration/aggregation and technology uptake by safety-critical systems/CPS in general. This can only really be properly realized if this team provides an ever-present pivot for CPS projects and CSAs, developing the support environment required to manage the multi-dimensional challenges.

They would have two support roles: the development and investigation of concepts that are provided by projects and programmes that had been identified as useful support assets (to the projects), but their implementation being normally unattainable in the scope of the projects and programmes.

On the development side, support to programmes would include, as an example, enhanced tool techniques for directed communication (the right information, to the right people, at the right time – especially for start-ups). The research instrument team provides prototyping tools to relevant CSAs that would support deployment in technology orchestration projects, who then test and further develop the tools. Support for the creation and testing of tools largely depends on results from the investigatory side. Some examples include:

- Inter-community supports like wiki-type project glossaries to manage the multiple perspectives (e.g. mediation between safety/security, medical/railway, SMEs/large enterprises).
- Multi-community access like digital passports, allowing users to access and test many research tools with the same account.
- Improved techniques like supporting management of intellectual property rights.
- Connecting contributions, such as a holistic view of open-source tool advancement across projects.

The investigatory side considers and proposes enhancements from the product-side perspective, for projects, programmes and industrial policy. These would be potential assets for promoting in particular aggregative technology uptake and longer-term profitability. Investigations would consider enhancements outside our normal fields of operation. Potential concepts include:

- Supporting the project environment for capitalization on and continuation of knowledge from multi-stakeholder interactions. Approaches for iterative improvement. Incentives, performance measures, mentoring.
- How are CPS-specific and aggregative technologies advancing, what is the funding flow to the contributing communities? Studies on benefits – but also consequences of lack of funding.
- Managed contributions, e.g. open-source results – rather than a default expectation, should be with respect to conditions (such as business model, maintenance, community building).
- Considerations for adapting the destination (industrial processes) to the new technologies; how to lift constraints at the product-side.
- Lighthouse initiatives within programmes (advancing structuring and management policies) may provide ideas to be explored.
- Currently, technology readiness levels (TRLs) measure the advancement of individual components rather than aggregations of components. A complementary approach, let us say aggregative-TRLs, is therefore required. This is not to be confused with the "integration readiness levels" measuring the interface between technologies (how they connect), rather than aggregating technologies (managing their combined effect).
- Supporting the development of a body of knowledge and teach the science of CPS engineering.
- Balancing local/national/European interests across networks. For instance, cross-border Digital Innovation Hubs (DIH) could complement the specific interests of regional or national DIHs.
- Policy on protection of EU business data (~B2B GDPR). CPS representation would be relevant here to consider the effects of such a policy on CPS technology advancement.
- Studies to advise/encourage industry towards longer-term strategies. This may also include changes in government regulation to shift from short-term competition of yearly quotas towards longer-term and more profitable competition and managing incentives where average employee turnaround is 3-4 years. No CPS-specific studies on corporate evolution seem to exist yet.

The proposed ways forward through this higher-level support from a CSA and a research instrument not only enables significant advancements for future large-scale safety-critical systems research, but also addresses the recommendations made by previous projects for CPS technology (such as Platforms4CPS [34], which represents an update of several roadmaps). These earlier recommendations included: collaboration and defragmentation of siloes; public understanding of the importance of CPS; supervisory support to draw together a common body of knowledge; and developing talent in order to maintain Europe's leadership and sovereignty of diverse technology aggregations for multi-domain applications including transport, manufacturing and health.

Acknowledgements

Contributors for the five groupings of communities

- **Rajendra Akerkar** is Professor and Head of Big Data Technologies at Western Norway Research Institute.
- **Djamila Aouada**, is Professor and Head of CVI2 Group, SnT Faculty, Université du Luxembourg.
- **Alessandra Bagnato** is Research Scientist and Head of Modelio Research at Softeam (Docaposte Group).
- **Miklós Györffi** is Senior European Affairs Analyst at the Eötvös Loránd Research Network and former staff member of the European Parliament.
- **Michael Henshaw** is Professor and Programme Director in Systems Engineering, Associate Dean for Teaching, Loughborough University.
- **Peter Gorm Larsen** is Professor and Head of the DIGIT Centre, Department of Electrical and Computer Engineering, Aarhus University.
- **Carles Hernandez Luz** is Senior Researcher in Processor Designs for Safety-Critical Systems, Universitat Politècnica de València
- **Hugo Daniel Macedo** is Researcher in the DIGIT Centre, Department of Electrical and Computer Engineering, Aarhus University.
- **Claudio Pastrone** is Head of IoT and Pervasive Technologies Research Area in LINKS Foundation.
- **Peter Popov** is Associate Dean (International), School of Mathematics, Computer Science and Engineering, City University London.
- **Claudio Sassanelli** is Professor Department of Mechanics, Mathematics and Management, Politecnico di Bari.
- **Marcus Völp**, is Professor and Head of CritiX Group, SnT Faculty, Université du Luxembourg.
- **Thorsten Weyer** is a Professor of software engineering, Technical Hochschule Mittelhessen (THM), Giessen.

References

[1] S. Paul, “D3.4.4. MERgE - Recommendations for Security and Safety Co-engineering v3 partA,” 22 April 2016. [Online]. Available: https://itea4.org/project/merge.html. [Accessed 2 December 2022].

[2] “RITICS: Research Institute in Trustworthy Interconnected Cyber-Physical Systems,” [Online]. Available: https://ritics.org/. [Accessed 2 December 2022].

[3] “ICRI_SAVe,” Intel, [Online]. Available: http://icri-save.de/the-institute/. [Accessed 2 December 2022].

[4] “IET Code of Practice: Cyber Security and Safety,” Institution of Engineering and Technology, [Online]. Available: https://electrical.theiet.org/guidance-codes-of-practice/publications-by-category/cyber-security/code-of-practice-cyber-security-and-safety/. [Accessed 2 December 2022].

[5] “MERgE: Multi-Concerns Interactions System Engineering,” ITEA 4, [Online]. Available: https://itea4.org/project/merge.html. [Accessed 2 December 2022].

[6] “SESAMO: Security and Safety Modelling on CORDIS,” European Commission, [Online]. Available: https://cordis.europa.eu/project/id/295354. [Accessed 2 December 2022].

[7] “AQUAS: Aggregated Quality Assurance for Systems,” [Online]. Available: https://aquas-project.eu/. [Accessed 2 December 2022].

[8] “CPS Swarm,” [Online]. Available: https://www.cpswarm.eu/. [Accessed 2 December 2022].

[9] SAE International, “Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles J3016_202104,” 30 04 2021. [Online]. Available: https://www.sae.org/standards/content/j3016_202104/.

[10] International Organization for Standardization,, “ISO 21448:2022 Road vehicles – Safety of the intended functionality,” 2022. [Online]. Available: https://www.iso.org/standard/77490.html.

[11] M. Roberti, “What Is the Internet of Things?,” RFID Journal, 2016. [Online]. Available: https://www.rfidjournal.com/question/what-is-the-internet-of-things. [Accessed 4 December 2022].

[12] “Internet of Things 08: International Conference for Industry and Academia,” Zurich, 26-28 March 2008. [Online]. Available: https://iot-conference.org/iot2008/prg/program.pdf.

[13] H.Boyes et al, “The industrial internet of things (IIoT): An analysis framework,” Computers in Industry, vol. 101, pp. 1-12, https://www.sciencedirect.com/science/article/pii/S0166361517307285, 2018.

[14] K.H. Wöhnert et al., “Secure Cyber-Physical Object Identification in Industrial IoT-Systems,” Procedia Manufacturing, vol. Volume 51, pp. 1221-1228, https://doi.org/10.1016/j.promfg.2020.10.171, 2020.

[15] S. Sachdev et al., “Voice-Controlled Autonomous Vehicle Using IoT,” Procedia Computer Science, vol. Volume 160, pp. 712-717, https://doi.org/10.1016/j.procs.2019.11.022, 2019.

[16] A.J. Ferrer et al., “Admission Control for Ad-hoc Edge Cloud,” Future Generation Computer Systems, pp. 548-562, https://doi.org/10.1016/j.future.2020.08.024, 2021.

[17] E.M. Jakobsen et al, “NGIoT D4.6: Report on IoT business model innovation patterns and acceleration support activities,” 30 September 2022. [Online]. Available: https://www.ngiot.eu/wp-content/uploads/sites/73/2022/09/EU-IoT_D4.6-IoT-BM-patterns-and-acceleration-support.pdf. [Accessed 4 December 2022].

[18] “SAFEXPLAIN: Safe and explainable critical embedded systems based on AI,” CORDIS, European Commission, [Online]. Available: https://cordis.europa.eu/project/id/101069595. [Accessed 4 December 2022].

[19] M. Wooldridge, An Introduction to MultiAgent Systems, 2nd Edition, Hoboken, New Jersey: Wiley, May 2009.

[20] C. Cares et al, “Agent-Oriented Engineering for Cyber-Physical Systems,” ICITS 2019: Information Technology and Systems in Advances in Intelligent Systems and Computing, vol. 918, pp. https://doi.org/10.1007/978-3-030-11890-7_10.

[21] “High Performance Computing,” European Commission, [Online]. Available: https://digital-strategy.ec.europa.eu/en/policies/high-performance-computing. [Accessed 4 December 2022].

[22] N. Rajovic et al, “Supercomputing with commodity CPUs: are mobile SoCs ready for HPC?,” in SC13: International Conference for High Performance Computing, Networking, Storage and Analysis, Denver, Colorado, 2013, https://dl.acm.org/doi/10.1145/2503210.2503281.

[23] “ Probabilistic real-time control of mixed-criticality multicore and manycore systems (PROXIMA),” [Online]. Available: http://proxima-project.eu/. [Accessed 4 December 2022].

[24] “MASTECS: Multicore Analysis Service and Tools for Embedded Critical Systems,” [Online]. Available: https://mastecs-project.eu/. [Accessed 4 December 2022].

[25] “Smart Anything Everywhere,” European Commission, [Online]. Available: https://smartanythingeverywhere.eu/. [Accessed 4 December 2022].

[26] M. W. Maier, “Architecting principles for systems-of-systems,” Systems Engineering, vol. 1, no. 4, pp. 267-284, 1998.

[27] “ISO/IEC/IEEE 15288:2015 Systems and software engineering – System life cycle processes,” May 2015. [Online]. Available: https://www.iso.org/standard/63711.html. [Accessed 4 December 2022].

[28] J. Dahmann, “1.4.3 System of Systems Pain Points,” in INCOSE International Symposium, Las Vegas, Nevada, 2014, https://doi.org/10.1002/j.2334-5837.2014.tb03138.x.

[29] S-V Rehm et al, “The Metaverse as Mediator between Technology, Trends, and the Digital Transformation of Society and Business,” Journal of Virtual Worlds Research, vol. 8, no. 2, pp. 1-6, https://doi.org/10.4101/jvwr.v8i2.7149 , 2015.

[30] Z. Sun et al, “Augmented tactile-perception and haptic-feedback rings as human-machine interfaces aiming for immersive interactions,” Nature Communications, vol. 13, pp. 5224, https://doi.org/10.1038/s41467-022-32745-8, 2022.

[31] S. Braun et al, “Requirements on Evolution Management of Product Lines in Automation Engineering,” IFAC Proceedings Volumes, vol. 45, no. 2, pp. 340-345, https://doi.org/10.3182/20120215-3-AT-3016.00060., 2012.

[32] M.M. Herterich, “The Impact of Cyber-physical Systems on Industrial Services in Manufacturing,” in 7th Industrial Product-Service Systems Conference, Saint-Étienne, 2015.

[33] P.G. Larsen et al, “A Cloud-Based Collaboration Platform for Model-Based Design of Cyber-Physical Systems,” in SIMULTECH 20: 10th International Conference on Simulation and Modeling Methodologies, Technologies and Applications, Online, 2020, https://www.scitepress.org/Link.aspx?doi=10.5220/0009892802630270.

[34] “Platforms4CPS,” [Online]. Available: https://www.platforms4cps.eu/. [Accessed 4 December 2022].

[35] M. J. d. C. Henshaw, “Systems of Systems, Cyber-Physical Systems, The Internet-of-Things…Whatever Next?,” INSIGHT, vol. 19, no. 3, pp. 51-54, 2016.

Charles Robinson is the research projects leader in critical embedded systems at Thales Research and Technology, France

This document is part of the HiPEAC Vision available at hipeac.net/vision.

This is release v.1, January 2023.

Cite as: C. Robinson. Bridging the stakeholder communities that produce cyber-physical systems. In M. Duranton et al., editors, HiPEAC Vision 2023, pages 32-43, Jan 2023.

DOI: 10.5281/zenodo.7462013

The HiPEAC project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement number 871174.