The SafeCOP ECSEL Project: Safe Cooperating Cyber-Physical Systems Using Wireless Communication

Pop, Paul; Scholle, Detlef; Hansson, Hans; Widforss, Gunnar; Rosqvist, Malin

Published in:
Proceedings of Euromicro Conference on Digital System Design 2016

Link to article, DOI:
[10.1109/DSD.2016.25](https://doi.org/10.1109/DSD.2016.25)

Publication date:
2016

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Pop, P., Scholle, D., Hansson, H., Widforss, G., & Rosqvist, M. (2016). The SafeCOP ECSEL Project: Safe Cooperating Cyber-Physical Systems Using Wireless Communication. In *Proceedings of Euromicro Conference on Digital System Design 2016* (pp. 532-538). IEEE. <https://doi.org/10.1109/DSD.2016.25>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

The SafeCOP ECSEL Project

Safe Cooperating Cyber-Physical Systems using Wireless Communication

Paul Pop
DTU Compute Dept.
Technical University of Denmark
Kgs. Lyngby, Denmark
paupo@dtu.dk

Detlef Scholle
Alten Sverige AB
Kista, Sweden
Detlef.Scholle@alten.se

Hans Hansson, Gunnar Widforss, Malin Rosqvist
School of Innovation, Design, and Engineering
Mälardalen University, Västerås, Sweden
{hans.hansson, gunnar.widforss,
malin.rosqvist}@mdh.se

Abstract—This paper presents an overview of the ECSEL project entitled “Safe Cooperating Cyber-Physical Systems using Wireless Communication” (SafeCOP), which runs during the period 2016–2019. SafeCOP targets safety-related Cooperating Cyber-Physical Systems (CO-CPS) characterised by use of wireless communication, multiple stakeholders, dynamic system definitions (openness), and unpredictable operating environments. SafeCOP will provide an approach to the safety assurance of CO-CPS, enabling thus their certification and development. The project will define a runtime manager architecture for runtime detection of abnormal behaviour, triggering if needed a safe degraded mode. SafeCOP will also develop methods and tools, which will be used to produce safety assurance evidence needed to certify cooperative functions. SafeCOP will extend current wireless technologies to ensure safe and secure cooperation. SafeCOP will also contribute to new standards and regulations, by providing certification authorities and standardization committees with the scientifically validated solutions needed to craft effective standards extended to also address cooperation and system-of-systems issues. The project has 28 partners from 6 European countries, and a budget of about 11 million Euros corresponding to about 1,300 person-months.

Keywords—cyber-physical systems; systems-of-systems; safety-assurance; wireless communication

I. INTRODUCTION

A *safety-critical system* is a system whose failure might endanger human life or the environment. When a system might harm humans or the environment (or is intended to mitigate or manage such harm), decision-makers require *pre-release safety assurance evidence* that it manages risk acceptably. In some domains, developers prepare an explicit *safety case* combining this evidence with safety arguments, whereas in other domains developers must show that their processes and work products conform to a relevant standard. For the purpose of this document, we call this safety evidence also a “safety case”, and the work in SafeCOP applies also for the domains which do not explicitly use safety cases. The conceptual basis for certification is that the pre-release (design-time) evidence anticipates the possible circumstances that can arise from the interaction between the system and the environment, to show that these interactions do not pose an unacceptable risk. Certification is

very expensive, and can add a development cost overhead of 25 to 100% [1] and in some cases even 10 times more. For example, “a commonly accepted rule of thumb is that development of safety-certified software costs roughly 10 times as much as non-certified software with equivalent functionality” [2].

If, after performing an initial *Hazard Analysis and Risk Assessment* (HARA), a system is deemed safety-related, it has to be certified. Certification is a “conformity of assessment” performed by a third party: a “certification authority”, e.g. an independent organization or a national authority. The certification process depends on the concrete application domain. However, the main ideas are common to all domains. The overall goal is to ensure freedom from unacceptable risk. Safety requirements typically consist of a functional part and an integrity level. A *Safety-Integrity Level* (SIL) captures the required level of risk reduction, and will dictate the development processes and certification procedures that have to be followed, depending on the standard, e.g., IEC 61508 [3], ISO 26262 [4], or RTCA DO-178B [5].

Once a system is certified, the safety certificate is typically valid only for a specific system configuration. This is the case, for example, in the avionics area, where the system is certified as a whole, and even small changes may result in a requirement for complete re-certification. The focus of recent research in safety assurance [6] has been to develop “modular” certification approaches. The idea is that a “modular safety certificate” can be given to an individual subsystem (module), and then these certificates will be manually composed into a system certificate. Thus, when a module is changed, the re-certification efforts can be isolated to the effects of that respective module. Some certification standards, such as IEC 61508 and ISO 26262, allow “modular safety cases” where the safety cases are composed. For example, ISO 26262 has the notion of a Safety-Element-out-of-Context. Recent projects, such as, RECOMP (Reduced Certification Costs Using Trusted Multicore Platforms, EU Artemis JU, 2010–2013) and SafeCer (Safety Certification of Software-Intensive Systems with Reusable Components, EU Artemis JU, 2011–2015), have proposed modular certification approaches, but these are not yet used in current practice.

II. COOPERATIVE OPEN CYBER-PHYSICAL SYSTEMS

SafeCOP addresses safety-related cooperating cyber-physical systems, characterised by use of **wireless** communication, multiple stakeholders, dynamic system definitions, and unpredictable operating environments. We refer to such systems-of-systems as **Cooperative Open Cyber-Physical Systems (CO-CPS)**. We assume that no single stakeholder has overall responsibility over the CO-CPS, that the cooperation relies on the wireless communication to perform a safety-relevant function, and that security issues are of concern. Following the taxonomy of Wilkies et al. [7], SafeCOP targets systems that are of the following three types: (i) use inter-system communication to reach a common goal; (ii) rely on communicated information from other systems in order to ensure safe and/or efficient operation; (iii) provide services that may compromise safety if the communication fails.

Such Cooperative Open Cyber-Physical Systems can successfully address several societal challenges. For example, cooperative vehicles (vehicle-to-vehicle, V2V, and vehicle-to-infrastructure, V2I) have been shown to reduce fuel consumption, reduce the number of accidents (including injuries and fatalities), result in productivity gains and congestion savings, resulting in annual savings of 1,300 to 2,000 billion Euros [8]. The US National Highway Traffic Safety Administration (NHTSA) has estimated that “*V2V safety applications have the potential to address approximately 80% of crashes for unimpaired drivers*”. For Road Weather Stations, which is a use case in the project, employing (V2I) and infrastructure-to-vehicle (I2V) communication modes, we can deliver critical up-to-date real-time road-weather data, which can increase traffic safety. In the maritime area, cooperative boats [9] can dramatically increase navigation safety, since, according to the IMO (International Maritime Organization), 75% of ship accidents worldwide are due to human error. CO-CPS can also be employed in the healthcare market, which is characterized by dramatically increasing costs. For example, cooperative robots can be used to reduce the amount of physical labour in hospitals. The use cases (UC) addressed in the project are summarized in Figure 1.

The development CO-CPS poses challenges that are not adequately addressed by existing practices nor standards. While careful safety-aware design and thorough safety assurance is required, no single manufacturer has design authority over or responsibility for the safety of a set of cooperative embedded systems. Developing a safety critical system typically requires making design decisions that trade-off safety concerns, functionality, cost, and other considerations. Achieving adequately safe cooperative cyber-physical systems requires arriving at, realising, and assuring a safe design even though participants in the design process are competitors reluctant to share all of their concerns or intricacies of designs with each other. Moreover, due to the cooperative and openness nature, many circumstances which have to be covered by the pre-release safety assurance are difficult to anticipate at design time in the case of CO-CPS.

III. PROJECT OBJECTIVES

The concrete objectives of SafeCOP are:

- **Objective 1.** *Develop a safety-assurance framework for CO-CPS.* The primary objective of SafeCOP **is to** propose an approach to the safety assurance of CO-CPS which will facilitate their certification and market release. This will create new applications and market segments, successfully addressing societal challenges.
- **Objective 2.** *Develop a reference “Runtime Manager” architecture to support the engineering and certification of CO-CPS.* SafeCOP will define and develop a reference “Runtime Manager” (which extends the reference platforms in the targeted application areas) that detects at runtime abnormal behaviour, triggering if needed a safe degraded mode. The verification, validation and simulation methods and tools developed as part of Objective 1 will be used to produce, besides the safety assurance evidence needed to certify cooperative functions, also the conditions that need to be observed by the Runtime Manager to ensure safety.
- **Objective 3.** *Extend the current wireless protocols for safe and secure cooperation.* SafeCOP will evaluate the adequacy of standard wireless technologies for CO-CPS

UC1. Cooperative moving of empty hospital beds	UC2. Cooperative bathymetry w/ boat platoons	UC3. Vehicle control loss warning	UC4. Vehicles and roadside units interaction	UC5. V2I cooperation for traffic management

Two-robot autonomous bed mover to wheel ordinary hospital beds to a central cleaning facility.	(Semi-) autonomous boats forming a platoon cooperate to perform bathymetry measurements.	When a vehicle loses functionality affecting others, all vehicles and the road infrastructure are notified.	Road weather stations (RWS) collect weather measurements, including from other vehicles, and distribute them.	V2I for <i>traffic management</i> using position and speed data from vehicle-borne transmitters to optimise traffic.

Figure 1. Use cases addressed in the project

to be used in the target application areas, and will propose an application-level “safety layer” on top of existing protocols to ensure safe and secure cooperation such that CO-CPS can be certified.

- **Objective 4.** *Contribute to new standards and regulations.* An important objective of SafeCOP is to contribute to new standards and regulations, e.g., provide the certification authorities and standardization committees with the scientifically validated solutions they will need to craft effective standards which have been extended to address cooperation and system-of-systems issues.
- **Objective 5.** *Demonstrate the usefulness of SafeCOP concepts in target applications.* We take five real-world applications in several domains and build demonstrator systems which show how CO-CPS can have concrete utility across a broad range of real commercial applications.

IV. CONCEPT AND APPROACH

Figure 2 presents the SafeCOP safety assurance concept. The approach in SafeCOP is to *restrict the behaviour* of the cooperative safety function at *runtime*, such that the *design-time* safety assurance evidence, with additional *monitoring at runtime*, is able to guarantee the safety requirements. Such an approach may require changes to the certification standards,

hence the objective to contribute to new standards and regulations. This is managed thanks to the SafeCOP project partners who are *safety assessors* (DNV GL, Safety Integrity, DTI), as well as members of standards committees. Additionally, the project is strengthened by an external advisory board, comprising people with vast safety assurance and security-related expertise. They will make sure that the innovations developed in SafeCOP are grounded in current certification practice and are aligned with the current efforts in the technical committees tasked with extending the certification standards.

Figure 2 shows two CPS systems, A and B, developed respectively by two organizations. Note that the systems addressed in SafeCOP may consist of several (arbitrary) cooperating cyber-physical systems. Each system has a cooperative subsystem co-responsible for the cooperative safety function. The communication is wireless. SafeCOP will extend the current state-of-the-art wireless protocols by creating an application-level library and related API that acts as a “safety layer” on top of the existing protocols. If this API is used for the communication, we guarantee that the communication has “high integrity”, i.e., trust that the contents of messages are not corrupted either unintentionally or intentionally. This is needed because otherwise our cooperative function A cannot trust and therefore cannot use messages from the other system (B) to implement its safety function. We will reuse security results from other ARTEMIS projects such as DEWI and from the Cooperation Reference

Figure 2. The SafeCOP safety assurance concept

Technology Platform (CRTP), and our focus is on delivering a solution which is not susceptible to security threats, such as man-in-the-middle attacks. Security concerns are not covered in detail in current safety standards, potentially resulting in systems that are successfully certified according to relevant safety standards, but that still are open to security threats that may jeopardize safety.

Each organization prepares their safety assurance evidence at design time. In SafeCOP, the safety case prepared by organization A is done by composing the evidence of the cooperative subsystem A with the *public* evidence from the cooperative subsystem B (i.e., without exposing the Intellectual Property, IP, of organization B). This is similar to *modular certification* allowed by certification standards, e.g., the SEooC (Safety Element out of Context) in ISO 26262. The difference is that modular certification typically does not hide the IP and does not address runtime safety assurance. In the proposal, we also refer to this as *composing* safety cases, since the safety-case for the cooperative function will be based on the individual safety-cases for subsystems A and B.

The safety case relies on constraints that restrict the runtime behaviour of the safety function, in order to anticipate at design time the circumstances that can occur at runtime. These runtime constraints are derived from the safety requirements and are enforced by a *Runtime Manager* (RM), which monitors the cooperative safety function. The integrity of cooperative subsystem A is guaranteed by system A's *architecture* and safety mechanisms. The public safety evidence of the subsystem B is refined into a set of *demands* which have to be fulfilled by subsystem B in order to ensure its integrity. We call this a "conditional safety case": the safety requirements are guaranteed only if the demands are satisfied.

If the Runtime Manager detects abnormal behaviour, or if the demands are not satisfied, the cooperative safety functionality is disabled. SafeCOP will analyse the requirements of the use cases and propose constraints based on the definition of the cooperative safety functions, such that safety is guaranteed. The safety case needs to also consider the risk of requirement violations, and how to provide failsafe fallback mechanisms. Considered mechanisms, including safety under such scenarios, will introduce additional constraints on the systems involved.

The Runtime Manager is implemented as software that needs to run on the CO-CPS system to ensure safe cooperation at runtime. This RM has to be "separated" from the functionality of the device, such that lower-criticality functions do not affect the functioning of the high-criticality RM. Some platforms (e.g., AUTOSAR¹) provide such separation mechanisms, but some (such as ROS²) do not, so these separation mechanisms (similar in concept to a "virtual machine", but more lightweight) will also have to be developed. The RM has to know what to monitor; the "Verification and Validation" methods and tools will produce safety requirements that need to be monitored at runtime. Once the RM detects a safety violation, it will have to fallback to a "degraded mode". The functionality of the "degraded mode" will have to be developed in the demonstrators, and it is specific to the function of the respective demonstrator. It is challenging

to develop useful and failsafe fallback functions, since often it is not appropriate to just stop (assuming there is a "safe stop" function).

The Runtime Manager contains a *monitoring module* that performs *data acquisition* to collect safety-related data, which is then analysed and included in the safety assurance evidence. In this context, we say that we have a "living safety case", which is updated with information collected at runtime to increase confidence. In this context we also say that the safety case is *incremental*, and it may support *provisional certificates* allowing usage in limited scenarios (e.g., initially only in-the-lab use). Through evidence collected at runtime, provisional certificates can be upgraded to cover more general usage scenarios. The qualification of the Runtime Manager and the constraints are part of the design-time safety assurance evidence.

V. WORK PACKAGES

As the central goal in the project is to provide an approach for the safety assurance and engineering of cooperative open cyber-physical systems (CO-CPS), we have organized the project around a number of use cases that feed requirements and problems into the research work-packages (WPs). The proposed solutions are then demonstrated and evaluated to ensure the feasibility of the approach.

The project is organized in seven WPs:

- WP1. Requirements,
- WP2. Safety assurance framework for CO-CPS,
- WP3. Safe and secure wireless cooperation,
- WP4. Platform and tool support for safety assurance,
- WP5. Demonstrators and evaluation,
- WP6. Dissemination and exploitation,
- WP7. Management.

The research and innovation work in the project is defined by the requirements derived in WP1. These requirements are coming from the use cases, which are part of the WP5 on demonstrators and evaluation. In WP1 we do the collection, refinement and consolidation of the requirements and research drivers. These consolidated requirements and research questions are then the basis of the work of WP2, WP3 and WP4, which provide solutions that—in turn—are applied and evaluated in several demonstrators in WP5. WP2 is concerned with the development of the SafeCOP safety assurance framework targeting CO-CPS, which communicate wirelessly. The wireless communication is addressed in WP3 where the goal is to extend the current protocols such that they provide the required levels of safety and security for CO-CPS. The safety assurance framework is supported by the SafeCOP platform, which consists of reference architecture for CO-CPS and methods and tools for producing safety assurance evidence. The work in WP2, WP3 and WP4 is performed in parallel, with interactions on the issues related to the assurance framework, wireless protocol, architecture and methods and tools. All WPs are structured such that intermediate

¹ <http://www.autosar.org>

² <http://www.ros.org>

evaluation of the approach is possible every year. WP7 is responsible for ensuring and monitoring the collaboration as described above; all WP leaders are part of this task. The dissemination of the major findings of the project will be done in WP6, and the management activities in WP7. A description of each WP follows.

WP1 Requirements. The goal of this work-package is to establish both the business case of the approach as well as the requirements for the solutions as they apply to different domains. The business cases establish goals that have to be achieved, and that can be assessed in the evaluation work-package (WP5). The requirements provide the specific constraints and problems that have to be solved in work packages WP2, WP3, WP4. The focus of WP1 is on safety and security related requirements for the implementation of cooperative safety functions required by the use cases.

WP2 Safety assurance framework for CO-CPS. The main objective of this work package is to develop a practical safety assurance framework for CO-CPS. After an evaluation of the state of the art on safety assurance, the WP proposes an assurance framework that can address the challenges of CO-CPS by combining pre-release safety assurance with runtime monitoring. The basis for this framework is a composable safety case, which contains “demands” placed on cooperative subsystems in order to provide safety “guarantees” for the cooperative safety function. In this WP we also evaluate and extend a safety analysis method called STAMP, which is suitable for systems with a lot of interactions. This work package also produces a set of scientifically proven recommendations for the certification of CO-CPS.

WP3 Safe and secure wireless cooperation. In SafeCOP we address cooperative open systems that communicate using wireless technologies. We are interested to elevate the state-of-practice to develop technologies that are both safe and secure, to be used in the context of CO-CPS. Hence, we start by evaluating standard wireless technologies that can potentially be used for cooperative safety functions, and we extend these wireless technologies to ensure and facilitate assurance of safety and security in cooperative embedded systems. Once a safe and secure communication solution is available, in this work package we are also interested to design distributed cooperation algorithms with safety-critical requirements.

WP4 Platform and tool support for safety assurance. The goal of this work package is to provide a platform and tool support for safety assurance. The WP defines a reference platform (hardware, OS and middleware), with the aim to guarantee the integrity of the cooperative function. We will extend the major platforms from each application area, e.g., AUTOSAR for automotive and ROS (which lacks safety mechanisms) for mobile robots. The novel component in the platform is a Runtime Manager, which enforces the cooperative function safety requirements providing a failsafe state in case of failure. This work package also extends the ARTEMIS Reference Tools Platform, with a focus on extending tool flows to support in efficiently producing safety evidence for certification.

WP5 Demonstrators and Evaluation. The consortium develops a number of demonstrators to show the applicability of the approach in different industrial areas. The demonstrators are built using the wireless technologies, platforms, methods and tools in WP3 and WP4, by applying the safety assurance process developed in WP2. This work-package also evaluates the results of the demonstrators. It provides various requirements input to WP1, and provide feedback that can be used to guide further research and development work in work packages WP2, WP3, and WP4.

WP6 Dissemination and exploitation. All partners advertise SafeCOP to their networks; academic, industrial, business or general public. This work-package includes setting up the project web site, producing newsletters, organization of workshops, demo booths, etc. An important component is also to liaison with standardization organizations to provide information about the results of the project. The objective of the exploitation phase is to identify and implement the actions necessary to maximize the market value, the business potential and the social benefits for the European Union of the project outcomes. The phase will be carried out using the consortium’s networks and other channels to explore vertical applications, use cases and disseminate commercially the solutions developed within the project. The exploitation will also address the standardization activities: the definition of new standards for safety requirements and the specification of methodologies for testing and compliance to the SafeCOP concept, will represent an important achievement/highlight of the project.

WP7 Management. This work package contains all tasks related to the management of the project, i.e. monitoring and reporting. Central to the success of the project will be the establishment of a good quality plan, risk management plan and communication plans to ensure good information flow between the partners. Moreover, this work package also includes knowledge and Intellectual Property Rights (IPR) management in the project.

VI. CONSORTIUM

The consortium is industry-led, consisting of 7 Large Enterprises, 10 Small and Medium Enterprises (SMEs), working with 6 universities and 5 Research Transfer Organisations. The partners are positioned across the full value chain, from technology providers, to system integrators, OEMs and end-users. The presence of 3 safety assessors and 6 members of standardization bodies facilitates the exploitation of safety assurance results. As already mentioned previously, beside the stakeholders represented by ECSEL JU monitoring the project, the project established an advisory board.

Special emphasis is taken on a balance between technology users and providers on the one side, and large companies, SMEs and researchers on the other. This balance will facilitate the technology transfer from theory into industrial practice. Particular emphasis has been put on the integration of SMEs. This can be seen on the quality and number of SMEs involved in the project. The enterprises (SMEs and LEs) include Original Equipment Manufacturers, system integrators, and end-users.

The project partners are from six European countries, with four representatives from the Nordic countries (Denmark, Finland, Norway and Sweden) and two from Southern Europe (Italy and Portugal). An overview of the number of partners per country can be found in Figure 3

The project coordinator is Alten Sverige, a Sweden-based LE with a presence in 20 countries. The LE's (Alten Sverige, DNVGL, GMVIS Skysoft, Intecs, Odense University Hospital, TEKEVER Autonomous Systems and Thales Italy) and SMEs (ALTE Visetec, Aitek, Impara, Intelligence Behind Things Solutions, Maritime Robotics, SITO, Qamcom Research & Technology, Ro Technology, Safety Integrity and Technicon) in the project cover several market domains with representatives from the automotive, maritime and healthcare sectors. Their presence ensures that the five use cases are properly grounded, that solutions are business-oriented, and that the final exploitation of the results reaches the right groups across multiple domains.

Having experience with nationally co-funded projects where a whole country has had to drop out, we have organised our five demonstrators in national units bound together by an international research “cap” with sufficient redundancy in expertise to cover the withdrawal of a single country if need be. Each of the national units are partners who have worked successfully together before, though not all on the same projects, and all of the university researchers involved have worked together previously in various subgroup combinations. All of the university partners and most of the industrial partners have previous experience with both national and international research projects, although three of the partner departments have not been involved in European projects before.

A. Setting up the consortium

The ARTEMIS and ECSEL funding instruments have promoted the assembly of very large and complex projects, often involving more than 100 person years, 8-100 partners and a (public) budget of between 0.4-42 million Euros. The average project has 25 partners and 9 million Euros total budget [10]. The strategy is to “think big” to gain “impact” and even if it is not primarily the size of the consortium that is meant, it is still an underlying message, that the larger it is, the more impact it will have. “The ARTEMIS mantra ‘think big’ does not mean that all projects have to be huge ones like the ARTEMIS CESAR project (Cost-efficient methods and processes for safety relevant embedded systems), which has about 58 partners and about €68 million of investment, it means thinking about the impact that the project will have” [11]. The dimension of the projects poses several challenges for its management. Hence it is not likely that all the staff from two partners ever meet in the project. The policy of promoting large and complex projects is also reflected in the support for proposal that is available in the ARTEMIS consortium-building events. At the breakout session all interested potential partners are welcome. There is no mechanism to allow the consortium leader to sort out undesired partners. The worst scenario is to walk off with 30-40 interested organisations, all of them expecting to be part of the proposal; limiting the consortium is a difficult task.

The funding of ARTEMIS/ECSEL is a blend of European contribution and contributions from each national innovation agency [12]. Each national agency has its own criteria and rules

Figure 3. The SafeCOP consortium

for payment. Most countries ask for an industrial project leader, and a specific budget ratio between industry and academia. That means that one prospective academic partner often has to find one or two other partners from the private sector to be nationally eligible. This means that the consortium will grow at least one extra round, without any real chance for the consortium leader to control the development.

Specific challenges are: the risk that large segments of partners or sub-clusters fall away, including valued partners, the risk that some sub-clusters cannot create eligible national consortia, and when some countries choose not to fund a specific project, or otherwise run short on budget—or, frankly, stop supporting the funding scheme.

A structured semi-open methodology. We have had experience with such a “snowball strategy” several times before. To take the lead and propose a topic and gather a consortium is not an easy task in a very open environment. As an alternative to the “snowball strategy”, we performed a more structured process in the SafeCOP proposal. That fosters narrower, smaller and (we believe) better consortia. The objectives for this are to gather a large group of interested potential partners, but through the process select the most desired ones.

We first proposed our SafeCOP project at a consortium-building event, early in February 2014. In this case we presented the project in a five-minute pitch talk, together with 50 other presenters in a plenum session. We also presented a poster, and

the project was also posted on the web a couple of weeks ahead. The result was a list of 37 interested individuals, representing 31 different organisations, where 4 were large companies or industries, 6 SMEs, 12 institutes and 9 universities, from 14 countries. The “usual” process would be to use the breakout sessions to form an initial outline of the proposal, and start assembling the consortium.

But for us, the next step was to contact the 37 person large group after two weeks. The message was that we planned to form a consortium out of the group of interested partners. They were all given the task to describe (1) their own organisation, (2) what their contribution would be and (3) whether they would be willing to lead any task. They were given a three-week deadline. The result was a detailed list of potential partners, but the list had been shortened to 10 potential partners, of whom 1 was from industry, 2 from SMEs, 3 from institutes and 4 from universities, from 10 countries. We believe that the action sorted out the better half of the list—those who were actually responsive to joint actions.

At the end of the day eligible country consortia are needed in this kind of call, and therefore the next step was to ask the 10 interested potential partners to provide national rules for the call (if known), and also to propose additional potential partners from their own country *if needed*, with respect both to national rules and the direction of the proposal. The potential partners had one week to suggest partners and another week to get the same kind of information from these new, suggested partners. At this stage at least one country left, but also one new entered. The result was a detailed list of potential partners, but the list had been extended to 26 potential partners, of whom 5 were from industry, 8 from SMEs, 7 from institutes and 6 from universities, from 10 countries.

Thereafter we selected three core partners, from three different countries (Denmark, Italy and Portugal); however, the Italian company could not commit itself at this stage. The core team worked out a “write-up” and selected partners and partner countries, mostly from the set of already interested partners, but also some totally new ones, that fitted into the project. Now the first revision of the consortium was Sweden, Denmark and Portugal, plus Norway, the Netherlands and Germany. In addition, Austria was asked to join. A message was issued for all the interested organisations that they were currently not included, but that they might be taken into account at a later stage. At this stage Italy re-entered into the consortium, while Austria, the Netherlands and Germany fell away.

We have established this way of working to find better ways to establish new European research consortia. First, we identify *la tête de la course*, as a core team, and then we pick the break-away specialist out of the bunch of the platoon—using a sports idiom. In this “marathon methodology” we try to select the best of those who want the most, to form a winning team.

VII. CONCLUSION

This paper has presented the SafeCOP ECSEL project. We have covered the societal challenges addressed, the project objectives, the overall approach and the consortium formation process. As mentioned, SafeCOP targets safety-related Cooperating

Cyber-Physical Systems (CO-CPS), where no single stakeholder has the overall responsibility over the resulted system-of-systems; safe cooperation relies on the wireless communication; and security is an important concern. Although such CO-CPS can successfully address several societal challenges, and can lead to new applications and new markets, their certification and development is not adequately addressed by existing practices. Note that many of the the research and innovations of SafeCOP also apply to CO-CPS that are not safety-related.

SafeCOP brings clear benefits in terms of cross-domain certification practice and implementations of cooperating systems in all addressed areas: healthcare, maritime, vehicle-to-vehicle and vehicle-to-infrastructure (V2I). The advantages include lower certification costs, increased trustworthiness of wireless communication, better management of increasing complexity, reduced effort for verification and validation, lower total system costs, shorter time to market and increased market share. The results are demonstrated in five demonstrators: cooperative moving of empty hospital beds, cooperative bathymetry with boat platoons, vehicle control loss warning, vehicle and roadside units’ interaction and V2I cooperation for traffic management.

ACKNOWLEDGMENT

The SafeCOP project is funded from the ECSEL Joint Undertaking under grant agreement n° 692529, and from National funding.

REFERENCES

- [1] IBM Software Rational, “DO-178B compliance: turn an overhead expense into a competitive advantage”, White paper, 2010.
- [2] K. Nilsen. “Certification requirements for safety-critical software.” RTC Magazine (2004).
- [3] IEC 61508: Functional safety of electrical/electronic/ programmable electronic safety-related systems. International Electrotechnical Commission, 2010.
- [4] ISO 26262 - Road vehicles—Functional safety. International Organization for Standardization / Technical Committee 22 (ISO/TC 22), 2009.
- [5] RTCA DO-178B. Software Considerations in Airborne Systems and Equipment Certification. Radio Technical Commission for Aeronautics (RTCA), 1992.
- [6] J. Rushby: “Modular Certification”, Technical report, NASA/CR-2002-212130, 2002.
- [7] T. Willke, P. Tientrakool and N. Maxemchuk, “A Survey of Inter-Vehicle Communication Protocols and Their Applications,” IEEE Communications Surveys & Tutorials, 11(2):3-20, 2009.
- [8] “Autonomous Cars: Self-Driving the New Auto Industry Paradigm”, Morgan Stanley report, November 6th, 2013.
- [9] I. Keddie. “Analysis: UMs set on course for naval service, Jane’s Defence Weekly”, 2012, HIS Global Ltd.
- [10] ARTEMIS-IA, <https://artemis-ia.eu/> 2015-02-13
- [11] The Co-summit 2013: Bolstering Software innovation to foster hi-tech employment and industry opportunities <https://itea3.org/press/the-co-summit-2013-bolstering-software-innova.html>, Published on 01 March 2014.
- [12] Decision of The Governing Board of The ECSEL Joint Undertaking. Adopting the ECSEL Work Plan for the year 2014, ECSEL-GB-2014.07, http://ecsel.eu/web/downloads/documents/ECSEL-GB-2014-07-workplan_v33.pdf 2015-02-13
- [13] E. Armengaud, “The CESAR Reference Technology Platform (RTP) – and way beyond”, presented at the ARTEMIS Spring Event, 13–14th of March, 2013, Brussels.