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

D7.3 Initial Exploitation Report

Version 1.0

Documentation Information

Contract Number	871669
Project Website	www.ampere-euproject.eu
Contractual Deadline	31.12.2020
Dissemination Level	PU
Nature	R
Author	Enkhtuvshin Janchivnyambuu (SYS)
Contributors	Zeeshan Ansar (SYS), Tommaso Cucinotta (SSSA), Scordino Claudio (EVI), Xavier Salazar (BSC), Frank K. Gürkaynak (ETHZ), Vergnaud Thomas (TRT), Luis Miguel Pinho (ISEP), Hamann Arne (BOS), MERLINI Marco (THALIT)
Reviewer	Frank K. Gürkaynak (ETHZ)
Keywords	Exploitation report, exploitable assets, intellectual property

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

Change Log

Version	Description Change
V0.1	Drafted first version
V0.2	Contributed by partners
V0.3	Revision based on comments from BSC
V1.0	Revision based on comments from ETHZ and THALIT

Table of Contents

Executive Summary

The European project AMPERE will develop a full ecosystem supporting and easing the development of future high-performance and real-time embedded applications that require the non-functional requirements (such as time predictability, energy-efficiency, safety and security) inherited from the cyber-physical interactions, on heterogeneous architecture including multi-core, GPU and FPGA acceleration.

Its objective to use the most advanced energy-efficient and highly-parallel heterogeneous platforms is to fully exploit the benefits of performance demanding emerging technologies, such as artificial intelligence or big data analytics. It will achieve that by the combination of model-driven engineering (MDE) and parallel execution, two important technical challenges at the system design and the computing software stack of CPS.

AMPERE project targets TRLs 2-5, and will deliver a working prototype tested in at least two different use cases described in the deliverable D1.1 “System models requirement and use case selection” (Intelligent Predictive Cruise Control and Obstacle Detection and Avoidance System) using the heterogenous platforms selected by partners in the deliverable D5.1 “Reference parallel heterogeneous hardware selection” (Xilinx UltraScale+ MPSoC and NVIDIA Jetson AGX), demonstrating improvement over the state of the art in industrial and professional domains.

This is the first exploitation report of AMPERE Work Package 7. An update of this document will be issued with the deliverable D7.5 Intermediate Exploitation Report in December 2021.

1. Introduction

Exploitation is a major commitment for AMPERE as proven by the diversity of industries in the consortium, and the IAB interested in the exploitation of AMPERE results. The exploitation activities will include the AMPERE ecosystem and the tools that form it and the definition of a roadmap for the adoption of parallel heterogeneous computing in safety critical systems. The fact that the AMPERE ecosystem is based on already existing tools (most of them owned by AMPERE partners) targeting COTS parallel heterogeneous computing platforms, provides two key exploitation advantages:

- the time-to-market will be reduced as most of these tools are already used in industry.
- it will enable partners to incorporate into their portfolio offerings tools supporting COTS low-energy parallel heterogeneous computing platforms, potentially increasing their revenues.

In the scope of the task 7.2 “Exploitation activities” of WP7, three exploitation reports will be delivered at each milestone as shown in Table 1. This deliverable, D7.3, describes our initial plans for the exploitation of the technologies and covers the first 12 months of activities carried out in the Task 7.2. This report will be updated as shown in Table 1.

Table 1:Exploitation reports

Deliverable	Deliverable Name	Deliverable Date
D7.3	Initial Exploitation Report	M12
D7.5	Intermediate Exploitation Report	M24
D7.7	Final Exploitation Report	M36

In this document we will describe our initial intentions for the exploitation of the technologies resulting from the project.

This deliverable presents an initial list of potential assets that will be made available as a result of the project. It gives an overview of management of knowledge and Intellectual Property Rights (IPR), and the different strategies for how to best utilize all exploitable assets. It also provides an overview of the current exploitation plans of each partner.

At this version of the document, joint exploitation plans by the project partners are not yet covered. We have planned to define such exploitations in the phase 4 of the project, and these will be included in the deliverable D7.7 ‘Final Exploitation Report’ in M36.

2. Management of Knowledge and Intellectual Property Rights (IPR)

The ownership and access to key AMPERE knowledge (IPR, data etc.) is an essential part of the Consortium Agreement (CA) signed by all partners. All IPR provisions will follow the spirit of the H2020 programme framework. The Section 3 “Rights and Obligations related To Background and Results” of CA defines the pre-existing partners’ know-how and describe the provision of its exploitation, the ownership of the project's results. Ownership of intellectual property shall be shared where there is joint invention. CA defines and regulates in detail the use of:

- **Background knowledge**. Information that can be relevant for the execution of the project, held by the partners prior to project start or acquired outside the project during the same project period. Access rights to this knowledge will be available to all partners only if they are valuable or useful

for carrying out project activities. Information may include (among others) the set of tools and software components integrated in the AMPERE ecosystem, i.e., model-driven approaches, parallel programming models, existing synthesis tools, compilation, operating systems and hypervisor, timing and schedulability analysis tools, energy models, resilient solutions, run-time parallel frameworks, offloading mechanisms, etc.

- **Foreground knowledge**. Results generated as an activity of the project, independently of whether they can be protected or not. Results may include (among others) new model-driven and parallel programming model extensions incorporating functional and non-functional information included in the meta-models, compiler techniques, multi-dimensional optimisation synthesis methods, timing analysis and scheduling techniques, energy efficiency methods, resilient solutions and functional safety mechanisms, operating system, run-time libraries, parallel programming models, applications, publications, reports, deliverables, roadmaps, etc. In general, all partners are bound by the terms and conditions of the Commission's contractual rules and a specific piece of foreground knowledge is the property of the partner(s) who has/have generated it. Each partner may use the results and material produced within the project for project purposes provided that such use does not come into conflict with the terms of the project Grant Agreement or the European legislation.
- **Patents**. In case a partner wants to submit a patent application, it must first inform the GA about its intention. The GA is in charge of handling any potential conflict prior to filing a patent. In case a conflict arises, all involved partners will notify the project coordinator that will guarantee that all partners are correctly represented and will guide the negotiation. Any conflicts will be addressed following the conflict resolution process presented in Section 3.2.3. Information of patent applications will be made available to the EU through regular management reports. The costs of the patent will be covered by the submitters.
- **Software/hardware accessories**. The software and hardware accessories from AMPERE partners (e.g. tools, components, devices, programs) required by other AMPERE partners to fulfil the project objectives shall only be used for the purpose of the project. Software products shall be made available free of charge, and hardware products at base costs including handling fees and depreciation. All these items shall be deleted or returned after the end of the project. These agreements shall be extended beyond the duration of project at partners' discretion.
- **Open Access**. The consortium is committed to provide at least green open access wherever feasible following the provisions of Horizon2020 guidelines. Green open-access is also known as self-archiving and means that authors deposit a preprint, a potentially revised author version or, where possible, a final peer-reviewed publisher's version of their publication at an institutional or subject repository that allows public access. Following open access policies of key publishers in our field (including SPRINGER, Elsevier, JSA, ACM, and IEEE) we have budgeted minor publication costs to allow for limited payments for open access.

3. Expected Exploitable Assets

AMPERE will devise a complete system design and computing software ecosystem including the stack for designing, implementing and efficiently executing dependable and physically-entangled systems on platforms composed of the most advanced COTS energy-efficient parallel heterogeneous architectures. With the objective of reducing the time-to-market and thus maximising exploitation opportunities, the AMPERE ecosystem is based on existing tools owned by AMPERE partners or IAB members (except GNU tools, which are open-source). Moreover, AMPERE will develop a powerful interface among the different tools to facilitate its integration on different development environments (supporting different parallel

heterogeneous platforms). This has been identified as a fundamental mechanism of the AMPERE ecosystem for successful exploitation after the end of the project.

Table 2 describes the different software layers and identifies the initial set of tools that we aim to incorporate into the AMPERE ecosystem, showing the corresponding owner (when known or relevant) and license of each tool. The AMPERE consortium will periodically update and assess the list of exploitable assets at each milestone using H2020 approach [1] that defines a Key Exploitable Result (KER) and add new details to list of the assets when possible. Next update to the list will take place in the beginning of the Phase 3 of the project as part of Deliverable D7.5, and after the phase 2 we expect to have a more refined view of all components of the AMPERE ecosystem.

Table 2 AMPERE System Design and Computing Software Ecosystem

Software Layer	**Description**	**Tool**	**License (owner)**
DSMLs	It will include enhanced model-driven languages capable of better expressing and verifying non-functional constraints such as performance, energy, safety and security, time predictability and fault tolerance in the context of the parallel heterogeneous computing, all aspects of notable importance to facilitate certifiability. The DSMLs will be used for describing the use-cases.	AUTOSAR	Proprietary (AUTOSAR)
		AMALTHEA	Open-source (BOS)
		CAPELLA	Proprietary (TRT)
Parallel programming models	The supported parallel programming model(s) will be determined by the selected parallel heterogeneous platforms upon which the use-cases will execute. In that respect, AMPERE will prioritize heterogeneous architectures supporting well-known programming models capable of distributing the computation across different platforms and expressing structured and unstructured parallelism with fine-grain support for synchronization, and featuring several acceleration devices, including many-cores, GPUs or FPGA.	OpenMP	Open-source (OpenMP)
		CUDA	Proprietary (NVIDIA)
		OpenCL	Open-source (Khronos)
		COMPSs	Open-source (BSC)
Artificial Intelligence	It will include deep learning frameworks and associated tools and APIs to generated optimized code for inference tasks in deep neural network executed on the selected parallel heterogeneous platforms.	TensorFlow	Open-source (Google)
Code synthesis tools	It will include novel code synthesis tools capable of transforming the DSML into an optimised parallel source code supported by the underlying parallel heterogeneous platform. This transformation will be optimized considering multiple optimization criteria.	Code synthesis tools	Open-source (AMPERE)
Analysis and testing tools	It will include a set of powerful analysis tools, capable of guaranteeing efficient multi-criteria optimizations at development phase, guiding the model-driven to programming model transformation and ensuring that functional and non-functional constraints are fulfilled. At deployment phase, these tools will be in charge of testing that the resultant parallel execution maintains the optimizations done at development phase.	Multi-criteria analysis	Open-source (AMPERE)
Compilers	It will include compilers capable of extracting the	Mercurium	Open-source

and hardware synthesis tools	control- and data-flow information needed by the analysis tools to maximise multi-criteria optimisations. Moreover, the compiler will support the selected parallel programming models to transform the parallel directives into the corresponding run-time calls. Finally, the ecosystem will include tools for compilation and synthesis of hardware blocks to be deployed on FPGA-enabled platforms.		(BSC)
		GCC/LLVM	Open-source
		Vivado	Proprietary (Xilinx)
Run-time libraries	It will include three different run-time libraries in charge of: (1) orchestrating the parallel execution as defined by the parallel programming model; (2) managing heterogeneous computing, including an efficient offloading of code and data to accelerator devices; and (3) supporting an efficient computation on accelerators, including DPR SoC-FPGAs. Moreover, these libraries will be extended with run-time monitoring and optimization techniques to: (1) track the parallel computation and ensure that non-functional requirements are constantly fulfilled, and (2) dynamically optimize the parallel execution for the multi-criteria constraints to better adapt the parallel execution to changes on workload conditions.	GOMP	Open-source (GNU-GCC)
		Nanos	Open-source (BSC)
		KMP	Open-source (LLVM)
		Vivado	Proprietary (Xilinx)
Operating systems	It will incorporate multiple operating systems (OSs) running on both host and accelerators. The range of the available OSs will include at least a general-purpose OS (Linux) as well as a certified Real-Time OS (RTOS) like AUTOSAR-compliant ERIKA Enterprise, the industrial-grade RTOS developed by partner Evidence. The selected operating systems will be properly enhanced to include additional functionalities that will allow to reduce the overall energy consumption while meeting the rest of non-functional constraints to support an efficient parallel execution.	Linux	Open-source
		ERIKA Enterprise	Open-source & commercial (EVI)
Hypervisors	It will incorporate a safe and secure real-time hypervisor capable of running on the target architecture and managing accesses to hardware resources. The hypervisor will execute applications and selected OSs with the needed safety assurance and provide secure separation of allocated resources and computational resources. The hypervisor will be enhanced with new functionalities to efficiently control processor resources and efficiently host run-time libraries.	PikeOS	Proprietary (SYS)

At the end of the AMPERE project implementation, the previously-identified KERs will be evaluated against the H2020's result maturity level matrix shown in Figure 1 to specify the stage of the project's results and readiness of technologies developed in the scope of AMPERE project. The revised KERs will be included in the deliverable D7.7 Final Exploitation Report.

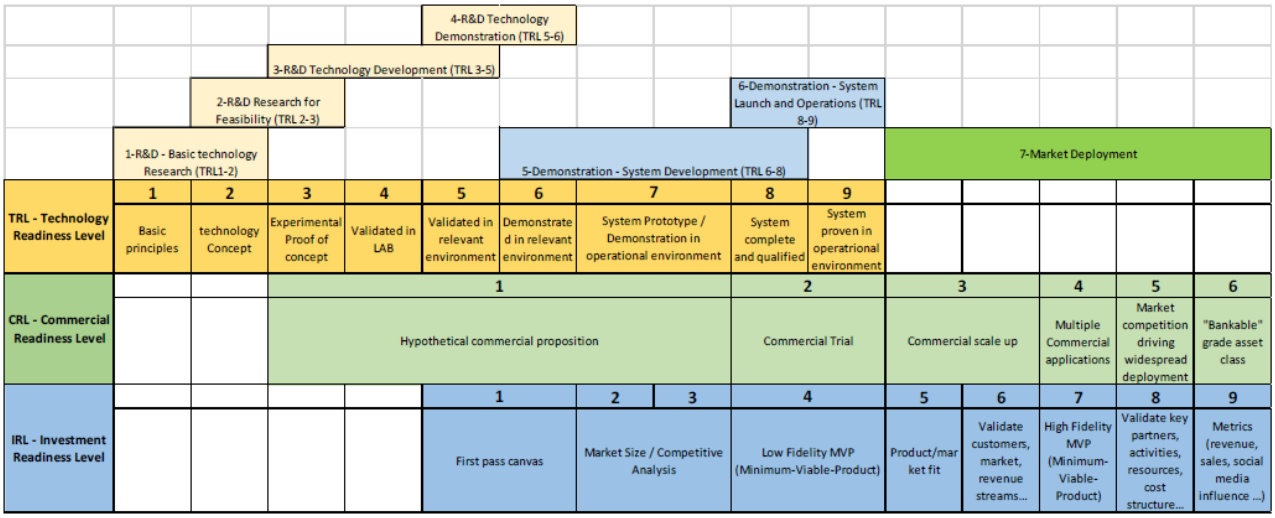

Figure 1: H2020's project result maturity level matrix

4. Exploitation Strategy

Due to the different business interests of the AMPERE project partners, several approaches will be used in order to make the project results more exploitable, to promote and to raise awareness and take-up of the project results. As for now the project is still in its early stages and without complete knowledge of what the final results of the project will be, we expect the following approaches to be the most relevant ones for AMPERE.

4.1. Licensing

The AMPERE ecosystem will be composed of tools (see Table 2 for the complete list) that either already exist (open-source and proprietary) or will be developed as needed, e.g. the synthesis tools. With respect to those developed from scratch, AMPERE intends to release them as open-source. Also, in task 5.2 an enhanced version of Linux to be run on the selected platform will be developed. All kernel source code developed within this task will be released through the popular platforms (e.g., GitHub) and the typical communication channels (e.g., Linux Kernel Mailing List and Linux conferences). However, decisions about licensing each components of the ecosystem will be made in the light of the progress and in alignment with the consortium agreement.

There are several advantages in adopting an open-source model, including effort and time-to-market reduction by re-using existing code, free help and support from a community of developers and potentially better code. In any case, it is important to understand the obligations and the dynamics of an open-source model, ensuring compliance with specific licenses and, at the same time, allowing effective exploitation activities that will not negatively affect the revenue streams for involved partners.

The AMPERE project will leverage the long experience of partner EVI, who has successfully adopted an open-source business model for more than a decade. In particular, the AMPERE consortium will carefully select the best license for each new open-source component developed in the project. For some components, the consortium will select industry-friendly licenses (i.e., BSD, MIT, LGPL, GPL with linking exception) that do not impose disclosure of source code of proprietary applications. For other novel components (where the business opportunities are higher) the consortium will evaluate a dual licensing approach [2]: a more restrictive license (e.g. GPL) allowing to spread the knowledge, and a commercial license (including additional services like e.g. technical support, bug fixing, updates, training and customization) allowing commercial exploitation of the project results. In case of components that were not developed within the AMPERE project – whose IP rights do not belong to the consortium and thus

cannot be released under a different license – the consortium will select only components that use industry-friendly open-source licenses. This way, the industrial partners will be able to use the AMPERE ecosystem for developing their own products without restrictions.

4.2. Integration of new software components and tools

Probably, one of the major challenges for an efficient exploitation of the AMPERE ecosystem is its capability to integrate new software components and tools to cope with the requirements that different application domains must face. It is therefore of paramount importance to facilitate the interoperability among the different tools to exploit AMPERE in multiple domains. In that respect, the primary goal of AMPERE will be to develop a powerful interface capable of gluing the different tools (based on XML or JSON) and its interaction to properly cover different development necessities. This has been identified as a fundamental mechanism of the AMPERE ecosystem for exploitation success after the end of the project. AMPERE includes a specific work package (WP6 - AMPERE System Design and Computing Software Ecosystem Integration) to address integration and interoperability issues among tools, and includes two different use cases (with different requirements) to ensure a proper assessment of tool integration features. Moreover, in the Task 5.4 the hypervisor will be enhanced to provide a Hardware Abstraction Layer (HAL) that will facilitate porting the AMPERE ecosystem to different COTS parallel heterogeneous computing platforms. To evaluate this feature, two different parallel platforms [3] have been selected for each use-case [4].

4.3. Standardization

The AMPERE project will provide recommendations to different standardization committees, to increase the synergies and collaboration between industrial and academic communities, and to align AMPERE to engineering practices and tools for current and future CPSoS development, allowing to efficiently uptake of the work after the end of the project.

4.3.1. Safety Standards

Functional safety refers to the property of the system to guarantee its correctness. It is verified by means of safety standards, which impose the code to be developed and analyse in a certain way (ISO26262 [5], IEC61508 [6], EN50126 [7], EN50128 [8] and EN50129 [9]). Those standards are of paramount importance to provide trust to system developers and end-users on the correct functionality of the underlying CPSoS controlling the product, in this case a car or a train. Unfortunately, the usage of parallel heterogeneous computing platforms in such systems is still challenging. Safety standards are a barrier to adopt parallel execution in general.

In that respect, the investigations that AMPERE will conduct regarding safety in the phase 3 of the project, will significantly advance towards the introduction of parallel heterogeneous computing into safety-critical systems, by evaluating the impact that model transformation and parallel execution have on functional and timing safety aspects.

To do so, the task 1.3 evaluates the functional safety aspects required at each integrity level defined in the safety standards ISO26262, IEC61508, EN50126, EN50128 and EN50129, which special interest on the isolation features required by each standard. Based on investigation results from this task, AMPERE intends to provide inputs to automotive and railway safety standards ISO26262 and IEC61508, EN50126 [7], EN50128 [8] and EN50129 [9], to which BOS and THALIT are members of the respective standardisation committee.

4.3.2. Software and System Architecture Standards

In the task 1.4 of the project, the DSML and parallel programming model will be extended to better capture functional and non-functional constraints, with the objective of deriving a more efficient parallel transformation, and then based on outcome of this task we also intend to propose recommendations to the AUTOSAR and OMG UCM standardization committees to better satisfy and secure composable features, and so enabling more efficient integration of new smarts functionalities into existing dependable and physically-entangled systems without compromising system security.

AUTOSAR DSML is regulated by an international committee, which provides development guidance and certification considerations to developers, facilitating the integration of components from multiple suppliers, and the verification of the correct functionality of the systems by means of isolation. Similarly, parallel programming specification committees, e.g. the IAB member OpenMP ARB or the OpenCL Khronos group, are in charge of overseeing the parallel programming model specification and produce and approve new versions of it to ensure an effective parallel computation within the new parallel and acceleration computing devices.

4.4. Open Access

The consortium is committed to provide at least green open access wherever feasible following the provisions of EU Horizon 2020 guidelines. Green open-access is also known as self-archiving and means that authors deposit a preprint, a potentially revised author version or, where possible, a final peer-reviewed publisher's version of their publication at an institutional or subject repository that allows public access. The project results disseminated as open data will also include the used benchmarks and evaluation results. This set of outcomes will provide sufficient evidence to appreciate the viability of the solutions in AMPERE and will enable, in subsequent research, to increase their Technology Readiness Level. Data will be shared with scientific peers to further the build-up on the AMPERE breakthroughs along the concepts of reproducible research.

Wide dissemination will be actively pursued by releasing all technical deliverables in the public domain and by posting the project results as open data in a public open-access repository.

In the task 8.3, the deliverable D8.3 Data Management Plan (DMP) [10] providing an analysis of the main elements of the data management policy that will be used by the applications with regard to all the datasets that will be generated by the project, has been prepared according to the Guidelines on Data Management in H2020 and delivered at month 6.

The deliverable DMP describes the data management life-cycle for all datasets to be collected, processed and/or generated along the lifetime of the project, and is applied to all datasets to be generated by the project.

4.5. Target Users

An important activity of AMPERE towards the dissemination and exploitation of its technology, will be to establish links to existent communities, targeting tools for system design, to promote, raise awareness and take-up of the project results. As a part of this activity, it is also important to identify early adopters and followers who directly use the technology developed in the scope of AMPERE at the first place.

The task to identify target users, especially early adopters and followers, will be carried out in the phase 3, and the required actions will be defined in the deliverable D7.5 Intermediate Exploitation Report as a part of the exploitation plan.

5. Current Market Analysis

Global economy and all industries are in turmoil as the coronavirus spreads across the world. As the situation continues to evolve, this section provides up-to-date market analysis.

The global embedded system market includes Tier 1 and 2 manufacturers such as NXP semiconductor, Infineon technologies, Intel corporation and STMicroelectronics which have their manufacturing facilities around the world. The COVID-19 not only affected their operations, but also affected their revenues as the companies using their products are also facing reduction in the consumer electronic devices. The decreasing sales in the automotive sectors, consumer devices, aerospace, and communication market is negatively impacting the semiconductor and embedded system market. While the healthcare industry is greatly growing due to extensive demand of ventilators and other advanced medical equipment across the globe. This is expected to positively impact the embedded system market in the overall market recovery.

According to the market research the embedded system market is estimated to be US$ 86.5 billion in 2020 and projected to reach US$ 116.5 billion by 2025 [11]. The main factors which is expected to fuel this growth in the coming five years are the current research and development activities in embedded system, rise in the demand of medical devices, smart energy devices, advanced driver-assistance systems (ADAS), portable devices, industrial automation and use of multicore processor and parallel heterogeneous computing platforms.

AMPERE is focusing on major embedded domains including but not limited to automotive, railway, industry automation and IoT by developing a software solution and analysis tools for the new heterogenous computing platform. The need of integrating heterogenous application is arising not only in the transport sectors such as automated driving and train control system, but also in other segments such as industrial and consumer electronics. AI and machine learning technologies are key to the current revolutionary era of industrial and smart devices. These AI-enabled smart devices require high processing power and real-time system support which can be provided by parallel and heterogenous computing platform. Our solution will contribute to the overall growth of the embedded system market especially impacting the European industry in automotive and I4.0.

In the automotive industry the solution for different car domains such as the cluster display, ADAS and ECU's needs to be meet constraints related real-time response for certification. The ability to perform the safety function is described in IEC61508 by Safety Integrity, which is a measure of the likelihood that a safety-relevant system will perform the specified safety functions under all specified conditions within a specified period of time. The highest safety integrity level (SIL) is defined by the hardware. The software inherits its SIL and must follow the processes specified in the corresponding software security standard. The IEC61508 has five levels SIL 0 to SIL 4, the ISO 26262 is based on four levels ASIL A to ASIL D (ASIL = Automotive SIL). By separating applications of different criticality, the cost of certification can also be significantly reduced in railway applications, as the entire system no longer needs to be certified at the highest required ASIL level. In addition, certified software components can be offered as COTS components and can be used in various projects without re-certification. For hardware, however, this approach does not work because it is usually constructed from different components that are currently inseparable. Nevertheless, the use of a common hardware platform leads to significant savings in procurement, development, and operation.

5.1. Market Situation: Automotive

According to the Statista research department [12], the worldwide car sales have decreased by over 17% in 2020 due to the coronavirus pandemic. It had been estimated pre-pandemic that the global car sales

will increase by 6% to reach its goal of 80 million sold cars. Due to the successful containment strategies, measures and government support, the economy in all major sectors is showing signs of recovery which is providing a lifeline for major manufactures in all part of the world. It is estimated that the global car sales will increase by 7% in the 2021. Similarly, the connected vehicle industry also faces decline in sales by 15% in 2020 which is expected to be recovered in 2021. ABI research expects 115 million global car sales by 2025 [12].

The trend towards higher levels of driving automation as well as the centralization of embedded Electrical/Electronic (E/E) architectures has led to a dramatic increase of complexity and required computation power for a lesser number of electronic control units. The resulting vehicle computers are one of the most demanding design endeavors in current cyber-physical systems industry. They are truly cyber-physical systems of systems (CPSoS) in the sense that multiple CPS which have previously been developed independently are now integrated and novel functionalities are developed on top of them in order to realize additional benefits for customers and society. Furthermore, these vehicle computers are developed by a multitude of partners supplying and integrating parts and functionalities on different levels and in several stages.

Thus, the goals of the AMPERE project are very relevant for the automotive industry and the deliverables, in-vehicle software components as well as the tool environment, are expected to play a significant role in developing vehicle computers. Due to the distributed development of vehicle computers, however, this requires that the underlying modeling and programming concepts will be widely adopted. Thus, a close alignment with existing DSMLs (AUTOSAR Classic and Adaptive) as well as a low entry financial barrier and vendor-independence are very important. For this purpose, the AMPERE ecosystem features the AMALTHEA system model, an open source system model adopted by many automotive players, as a central entry point for the automotive domain. Following this strategy, extensions and basic tooling will be made available as part of the Eclipse project APP4MC, which is the development environment around AMALTHEA.

5.2. Market Situation: Light Railway

Covid-19 also hit the railway industry globally. The revenue has decreased by 37.1% in 2020 according to Statista research department [13]. However, fortunately due to remarkable advancements in the preparation of vaccine and containment measures the industry is showing signs of recovery. The annual growth rate is predicted to increase by 16.8% (CAGR 2020-2025) to reach the projected market volume of US$149,387m by 2025 [13]. Annual market size of the rail control for Light Rail Systems is 1.1 Bn € and it is expected to grow 1.9% per annum in the future years.

The Light Rail Transit or tram market lacks global standardization and the segments are divided by geographic area and mode of use. Europe and APAC are using similar approach compared with Middle East or NAFTA area. Tram, LRT, Tram-Train or Tram-Metro may differ according to requested performances and physical constraints.

A major trend to “Autonomous Vehicle” also influences the railway world. However, public transport systems still rely on legacy technologies invented in the past century to ensure the safe movement of their users, especially in Signalling systems. The window of this market creates opportunities to innovative companies to introduce more smart systems to the industry. Europe is the area with more than 35000 Light Rail Vehicles in operation. Thus, for technology transformation in this sector the Europe would be the reference market as a showcase for a wider world market.

Self-localization and obstacle detection are an important and challenging issue in current driving assistance and autonomous driving research activities. Innovative methods for localization, awareness of surrounding environment and mapping will impact, in the next years, the railway world to introduce new technologies available today to Railway Signalling with application of them within the Safety Domain.

6. Individual Exploitation Plans

This section contains the individual exploitation plans by the project's partners. These plans express the initial works to be done to promote use of the technologies developed within the AMPERE project in the own respective domain after the completion of the AMPERE project.

6.1. Initial Exploitation Plans From Academic Partners

The exploitation plan and interests that are typically described by research partners are the following:

- Extend the capabilities of the system modelling and specification language technologies with additional features for parallel programming model and heterogeneous platforms.
- Research various system models (AUTOSAR, CAPELLA, AMALTHEA) to create new reference system design and computing software ecosystem for the future dependable and physically entangled systems with high-performance requirements.
- Investigate further techniques, and methods for the integration of parallel programming models and real-time analysis techniques, targeting more accurate and realistic analysis methods.
- Investigate reliable, robust and energy-aware solutions for autonomous and safety-critical systems.
- Evaluate engineering techniques supporting the design operation continuum of dynamic CPSoS and emerging technologies such as big data analytics or artificial intelligence.

The AMPERE project results should lead to new research by the partners and these future results will contribute to the continued evolution of the AMPERE model-driven framework.

6.1.1. Exploitation Plan of Partner (SSSA)

SSSA exploitation plans for the project concern mainly 3 areas:

- Hardware and software developers' communities
- Linux kernel community
- Academic real-time research community
- Automotive model-based community

Concerning hardware/software developers' communities, SSSA is planning to release the support of the FRED framework for Ultrascale+ on Linux as open source, together with its accessory hardware designs and demos. A ready-to-use package will be available to test the technology on a board equipped with the Ultrascale+ SoC. Xilinx will be also informed about the support in order to push for a wider adoption and magnify the impact of AMPERE in the development of next-generation CPSs that require FPGA-based hardware acceleration.

For what regards the Linux kernel community, our exploitation plans are to trigger fruitful discussions about the progress made within AMPERE on the side of energy-aware scheduling of real-time tasks on heterogeneous platforms. In this area, we plan to integrate the mechanisms to be investigated in AMPERE within the SCHED_DEADLINE code base, which is a mechanism originally developed jointly by SSSA and EVD, and still actively contributed to by both partners, and release any code as open-source patches. Therefore, we can operate alongside three main channels. First, we can submit possible patches to the SCHED_DEADLINE in-kernel scheduler on the Linux Kernel Mailing List (LKML) for review by the kernel community. Second, we can discuss the patch live at related events organized periodically. AMPERE project partners SSSA and EVD are both involved in the yearly organization of OSPM, the international workshop on energy efficient scheduling in the Linux kernel, which has been organized yearly in Pisa, with a strong industrial focus. Indeed, the workshop has been organized in close cooperation with, and attended by, major international industrial players and stakeholders in the area, including Intel, ARM,

Google, VMWare, IBM and RedHat. OSPM constitutes an excellent opportunity to gather feedback from the community of Linux kernel developers actively involved in the continuous improvement of the kernel features, especially on the side of real-time performance of applications. Third, we plan to investigate possible interest by users of the novel real-time mechanisms within the Linux kernel, leveraging the yearly Linux Plumbers Conference (LPC). This is a yearly meeting where many industrial practitioners gather, who are working on Linux from a 360-degrees perspective. This involves kernel main core developers, people working on device driver development, application and middleware developers, and Linux distributors and maintainers.

SSSA is finally planning to raise the awareness on the predictability of FPGA acceleration and the capabilities of the AMPERE ecosystem in the embedded/CPS and real-time research communities. To this purpose, we are evaluating the possibility of organizing a workshop on predictable heterogeneous computing within one of the major international events (e.g., the CPS Week, or the Embedded Systems Week). AMPERE partners will be invited to present their work on the AMPERE ecosystem, together with other presentations to be provided by other influential researchers working in the field. Also, some of the techniques investigated in AMPERE for energy-aware scheduling of real-time workloads, will be prototyped by modifying RTSim, an open-source simulation platform well-known among researchers in the community to simulate the timing of real-time applications. We plan to make our final scheduling algorithms for heterogeneous platforms available as an open-source tool that can easily be reused to improve the techniques and develop further research in the area.

For what regards the automotive model-based community, we plan to strongly interact with Bosch to have the modifications needed on the Amalthea models to support energy-efficient scheduling of real-time workloads on heterogeneous platforms, as well as the accompanying analysis and verification mechanisms to be realized, available as open-source patches to the App4MC open-source code base, or as independent open-source tools to be integrated within the App4MC design workflow.

6.1.2. Exploitation Plan of Partner (BSC)

The objective of the Barcelona Supercomputing Center (BSC) is to create a new reference system design and computing software ecosystem upon which the future dependable and physically entangled systems with high-performance requirements will be developed. To do so, BSC will integrate the outcome of AMPERE with the outcome of the FP7 project P-SOCRATES (addressing time predictability in OpenMP), and the H2020 projects AXIOM and LEGaTO (addressing energy and fault tolerance with OmpSs@FPGA), CLASS and ELASTIC (addressing real-time big-data analytics with COMPSs), and HP4S (addressing the use of OpenMP is payload systems for the space domain) industrial projects. In that respect, it is worth mentioning that BSC has already successfully developed two reference ecosystems widely used in the HPC, big-data and AI computing domains, i.e. OmpSs and COMPSs respectively. Moreover, the advances of OmpSs has deeply influenced the OpenMP standard. Finally, BSC will continue the research towards a complete integration of HPC and CPS computing domains by targeting further synthesis tools and extensions on well-known parallel programming models to better express functional and non-functional constraints of future dependable and physically entangled systems, enabling to apply the AMPERE technology to wider application domains, such as big data and AI applications, in which BSC is also very active.

6.1.3. Exploitation Plan of Partner (ETHZ)

ETH Zurich will continue research activities around energy-efficient heterogeneous many-cores, and the integration of such resource-management policies within well-established parallel programming models for such platforms. The increased knowledge in this area will:

- foster the education and training of experts in the technologies developed in the project, at all levels.

- position ETHZ as a strong industry-oriented applicative research partner.

Towards the end of the project, the developed technologies will have reached the adequate maturity for industrial uptake. The fact that such technologies will be evaluated in cooperation industrial partners (including discussions with IAB industrial partners) from the ground up strengthens ETHZ as a key partner for future technology transfer efforts. Furthermore, the commitment of AMPERE to interact with the EPI project and the use of HERO as a proto-typing platform enables ETHZ to extend its leading position in promoting of RISC-V, as part of the RISC-V foundation, where it is a founding member.

The maturity of power and energy management software and hardware layers developed within AMPERE will also enable ETHZ to push these techniques to broad audiences through its open-source PULP and HERO platforms, to maximize their impact in both industry and academia.

6.1.4. Exploitation Plan of Partner (ISEP)

ISEP will exploit project results through further research and development activities on the integration of parallel models and real-time analysis techniques, targeting more accurate and realistic analysis methods and fostering the use of the combined techniques in real-time embedded applications. ISEP will exploit these results with the companies in AMPERE, and in its IAB, addressing the transfer and take-up of results in the industrial domain, in cyber-physical systems in general, but with particular focus in rail and automotive. The outputs of ISEP in the project will be also integrated in the AMPERE framework, and exploited together with the project partners.

The knowledge obtained in the project will also be explored through publication of scientific results in international events and journals, and integration in the master and doctoral programs, where the ISEP team is involved.

6.2. Initial Exploitation Plans From Industrial Partners

The exploitation plan and interests that are typically described by industrial partners are the following:

- Productise the technologies developed within the AMPERE project to make them available as software tools and platform offers for the industry domains targeted by each partner
- Utilise AMPERE technologies for development of dependable and physically-entangled systems in their respective industrial domains
- Encourage their suppliers to utilise AMPERE technologies for the development of parallel critical real-time smart systems
- Establish distribution of AMPERE technologies in markets addressed by each partner
- Conduct promotional and marketing actions to create further awareness of the new AMPERE tools for addressing heterogeneous platforms and parallelised software development
- Undertake further development of AMPERE technologies as enhancements and updates to provide improvements and greater stability in the prototype technologies from the project
- Expand the portfolio of related products by collaborating with other organisations that can provide additional tools and platform technologies
- Support the take-up of AMPERE project technologies through awareness building amongst the European software development community.

In carrying out the above exploitation actions, all industrial partners will seek to create revenues and other commercial opportunities as important actors within the AMPERE project.

6.2.1. Exploitation Plan of Partner (EVI)

The work done by EVI in AMPERE will allow the company to enhance its AUTOSAR Classic solution based on the ERIKA Enterprise RTOS, being able of creating a competitive product for the European automotive

market. The main exploitable results consist on the support for both SYSGO's PikeOS hypervisor and RISC-V technology (expected to become more and more strategic in the forthcoming years). The exploitation plan for these parts will need to take into account potential licensing constraints imposed by the AUTOSAR consortium. The exploitation activity will then naturally flow through Evidence's shareholder Huawei, which could complement Evidence's offering with additional hardware and software components add-ons. The marketing activities will mostly consist on meetings with European automotive OEMs and Tier-1s. However, EVI will also collaborate with the other AMPERE partners to organize dissemination events (e.g. participation to conferences and exhibitions, organization of workshops, submission of papers, etc.).

In addition, within AMPERE EVI will investigate the development of a novel POSIX PSE53 RTOS targeting the automotive market. Such preliminary design is expected to grow through further private investments, becoming a potential AUTOSAR Adaptive solution. Once ready, a proper exploitation plan will be developed, which could consist either on a traditional royalty-based commercial offering or on an open-source business model. To this aim, a survey of both commercial and open-source solutions available will be carried out next to the end of the AMPERE project, to identify the best exploitation and dissemination strategies.

6.2.2. Exploitation Plan of Partner (SYS)

SYSGO exploitation plan as stated in the project proposal is to aim to strengthen position of PikeOS on European market which is currently dominated by US companies and create new markets for PikeOS. The strategy to achieve these goals is to constantly innovate in the product line. AMPERE will enable extending the PikeOS functionality to support parallel heterogeneous hardware. The targeted simplicity of programming of that heterogeneous hardware will enhance the exiting PikeOS market offering for dependability and real-time with cross-layer programming support to reduce programming and integration efforts and shift decision in favour of PikeOS. Existing customers plan and/or have to use highly heterogeneous hardware in next generation systems. Therefore, SYS sees support for parallel heterogeneous platforms as an important milestone. This support shall be easy to use and not introduce technical obstacles.

The exploitation plan is split into two phases during the project and after the project. During the project SYSGO will:

- create awareness of the problem and developed solutions by writing white papers and industry-oriented publications;
- keep main interested partners and customers updated on project intermediate results;
- show-case project result on the industrial trade shows focused on cyber-physical systems such as Embedded World, RTS, Avionics Europe, SAE, it-SA;
- closely interact with the product development team to decrease transition into product phase;
- investigate the exploitation opportunities with the use-case providers.

After the project, SYSGO will:

- publish press-release on new features co-located with trade-show (within 1 year); present results and benchmarks to interested customers and partners (within 1 year);
- integrate developed support and extensions for heterogeneous HW into certification base product (within 2-3 years);
- work with SYSGO’s multiplier partners such as design tool providers to integrate develop methodology into different tools (within 3-4 years).

6.2.3. Exploitation Plan of Partner (TRT)

One of the TRT's roles in the Thales group is to identify promising technologies and help Thales business units adopt them in order to improve the Thales industrial processes.

In project AMPERE, TRT brings its expertise in system modelling with CAPELLA. CAPELLA is a system modeller released in open source by Thales. It implements the Arcadia methodology to design system architectures. Compared with the other technologies involved in AMPERE, CAPELLA is at a higher level of abstraction. A large part of the AMPERE partners is working on lower levels (code generation for parallel computation, power consumption analysis, timing analysis, etc.). In the scope of AMPERE, TRT aims at gaining expertise in the performance analysis of architectures for high-performance computing. In particular, this covers the ability to address trade-offs between computation time and power consumption for parallel computing.

The intent of TRT is to gain knowledge in this field and apply it to Capella to enable the design of system architecture models that contain the necessary information to produce lower-level models for parallel computing architecture. Thus, TRT can assist Thales business unit in designing system architectures that target parallel computing using model-driven engineering (MDE). The MDE community in Thales gathers more than 500 engineers, across all Thales business units; a large part of them works with Capella.

In Ampere, TRT will directly assist THALIT to validate the feasibility of the Capella modelling approach. Beyond Ampere, TRT will be able to assist other Thales business units – first in the transportation domain, then in the whole Thales MDE community – to help them adopt the Ampere technologies in their design processes.

6.2.4. Exploitation Plan of Partner (THALIT)

Thales Italy (THALIT) has interest in testing and commercially exploiting the AMPERE technology for urban transport applications such as Light Rail Transit (LRT) and urban rail. By leveraging on the technological solutions developed in AMPERE, THALIT strives to improve the innovative autonomous tram solutions and building blocks delivered for the smart urban transport systems.

The Florence tramway receives 14 million passengers each year on its trams (Data 2017). The reasons that lead to the choice of the tramway are comfort, safety, economic convenience and certainty of travel times. The 87% of passengers are overall satisfied with the service and consider it good and excellent.
Due to geographical, infrastructural and socio-economic factors, the city of Florence plays a central role in the mobility scenario of the area, with about 80.000 and 70.000 cars moving respectively to and from Florence every day, accounting for severe traffic congestions and affecting quality of life.

To face these issues, the local governments (municipal, metropolitan and regional) have started a thorough transformation of the public transportation network, which includes the construction of new light rail lines, a complete redesign of the bus service and a widespread use of ICT solutions.

The development of the ODAS for the railway use case in the AMPERE project will contribute in:

1. Decrease the number of accidents in the tramway environment, warning the driver of a potential collision with detected target and reducing its reaction time;
2. Improve the state of the art, developing innovative technologies supporting compute-intensive applications in the railway and automotive domains;
3. Develop a meta model-driven abstraction, incorporating functional and non-functional constraints;
4. Reducing the development and integration costs of the ODAS functional critical software modules;

5. Improve object detection capability and reducing false alarms in critical environmental conditions (e.g. rain, fog presence, etc.);
6. Assess the impact of synthesis methods and parallel heterogeneous execution with regard to the safety and security requirements;
7. Extend parallel programming models to make a more efficient use of resilience redundancy techniques.

The key component of the AMPERE ecosystem to fully exploit the performance capabilities of platforms composed of parallel heterogeneous architectures will be the code synthesis tool, interfacing between the Capella DSML and the parallel programming model supported by the platform.

6.2.5. Exploitation Plan of Partner (BOS)

Regarding the prospects of economic success, no changes to the exploitation plan stated in the project description are necessary. However, it is expected that the use of the AMPERE ecosystem will cover a larger circle of users than originally thought. The possibilities of such a system are increasingly attracting interest inside Bosch. More precisely, in several discussions with large ongoing development projects within Bosch, the AMPERE technology gained a lot of interest despite being hindered by COVID-19 restrictions that limited working time.

The reason for that interest is that AMPERE directly addresses the challenges in model-centric development arising from the technology shift from dedicated µC-based to massively parallel and heterogeneous computing platforms. Especially the mapping methodology constructively guaranteeing timing properties has potential to master system complexity, and thus reduce development and validation efforts.

Bosch, therefore, intends, like stated in the project description, to introduce AMPERE methods within 2-3 years after project end into its internal development processes for systems with stringent timing requirements. Additionally, Bosch will actively and continuously participate in the standardization of the key results into industrial standards, especially AUTOSAR and ISO26262, and in already existing development environments such as AMALTHEA, included in the APP4MC Eclipse framework. First activities are already well advanced. For instance, Bosch extended the internal “Synthetic Load Generator” tool to support ROS2 middleware concepts (which are conceptually strongly related to AUTOSAR Adaptive) and made it available via open source for collaboration with partners (internal and external to the AMPERE project). Also, the open source APP4MC Eclipse framework (which is the basis for Bosch internal development across different business units) is currently extended to cope with AMPERE technology concerning the modelling of parallel execution on heterogeneous Systems-on-Chip. This “upstream first” strategy pursued by Bosch within AMPERE is crucial to strengthen the engineering backbone serving as basis for distributed development projects with many suppliers that are common in the automotive domain. This way the AMPERE technology will be broadly disseminated and exploited to its full potential in the automotive industry.

Moreover, the scientific potential and success of AMPERE are excellent. Conferences and other events where we contribute and publish AMPERE results are raising the profile of AMPERE, and are creating a wider circle of interest.

7. Planned actions

Each exploitation report will not be independent, but extension of previous exploitation report. For instance, D7.5 Intermediate exploitation report will be the updated version of D7.3.

Table 3 depicts the planned actions to be carried out as a part of Task 7.2. Outcome of each step will be reflected on respective deliverable.

Table 3. Planned actions

Project Phase	Action	Deliverable	Due date
Phase 3	Identify target users (early adopters and followers)	D7.5 Intermediate Exploitation Report	M24
Phase 3	Re-identify key exploitable assets (software components of the AMPERE ecosystem and others)	D7.5 Intermediate Exploitation Report	M24
Phase 3	Compatibility analysis of licenses for software components	D7.5 Intermediate Exploitation Report	M24
Phase 3	Revise individual plan for exploitation	D7.5 Intermediate Exploitation Report	M24
Phase 3	Evaluate the feasibility of exploiting a separate “bundle” of components	D7.5 Intermediate Exploitation Report	M24
Phase 4	Joint plan for exploitation	D7.7 Final Exploitation Report	M36
Phase 4	Results of the evaluation process	D7.7 Final Exploitation Report	M36
Phase 4	Define the procedures to protect IPR of individual tools	D7.7 Final Exploitation Report	M36
Phase 4	Recommendations to AUTOSAR and OMG UCM	D7.7 Final Exploitation Report	M36
Phase 4	Recommendations to safety standards	D7.7 Final Exploitation Report	M36

8. Acronyms and Abbreviations

- AI – Artificial Intelligence
- AMPERE – A Model-driven development framework for highly Parallel and EneRgy-Efficient computation supporting multi-criteria optimization
- API – Application Program Interface
- AUTOSAR – Automotive Open System Architecture
- COTS – Commercial Off-The-Shelf
- CPS – Cyber Physical System
- CPSoS – Cyber Physical System of Systems
- CPU – Central Processing Unit
- DMP – Data Management Plan
- DPR – Dynamic Partial Reconfiguration
- DSML – Directory Services Markup Language
- FPGA – Field Programmable Gate Array
- GPU – Graphics Processing Unit
- HAL – Hardware Abstraction Layer
- HERO – Heterogenous Research Platform
- IPR – Intellectual Property Rights
- KER – Key Exploitable Result
- MDE – Model-Driven Engineering
- MPSoC – Multiprocessor SoC
- OEM – Original Equipment Manufacturer
- OS – Operating System
- PULP – Parallel Ultra-Low Power
- RTOS – Real-Time Operating System
- SoC – System-on-Chip

9. References

[1] https://webgate.ec.europa.eu/funding/display/ECResearchGMS/Projects+and+results

[2] Multi-licensing, https://en.wikipedia.org/wiki/Multi-licensing

[3] Deliverable D5.1 Reference parallel heterogeneous hardware selection

[4] D1.1 System models requirement and use case selection

[5] ISO 26262-1:2011, Road vehicles - Functional safety

[6] IEC 61508 - Functional Safety, http://www.iec.ch/functionalsafety/

[7] EN 50126 - Railway Applications, https://standards.globalspec.com/std/260302/EN%2050126

[8] EN 50128 - Railway Applications, https://standards.globalspec.com/std/14317747/EN%2050128

[9] EN 50129 - Railway Applications, https://standards.globalspec.com/std/13113133/EN%2050129

[10] Deliverable D8.3 Data Management Plan (DMP)

[11] https://icrowdnewswire.com/2020/10/14/covid-19-impact-on-the-global-embedded-system-market/

[12] https://www.statista.com/topics/1487/automotive-industry/

[13] https://www.statista.com/outlook/271/100/trains/worldwide