


# OpenMP


## SC23 Booth Talk Series



## Exotic OpenMP Use Cases

Eduardo Quiñones, BSC

# Heterogeneous and Parallel Computing



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten  
New plot and data collected for 2010-2017 by K. Rupp

**Heterogeneous and Parallel computing** becomes key to cope with performance requirements



Network of HW/SW components that **must** operate **correctly** in response to its inputs from both **functional** and **non-functional** perspectives

Massively parallel systems that operate **as fast as possible**



Genomics



Weather



Big data

## HPC Domain (~300W)



**NVIDIA A100**  
(GPU-based)



**AMD Instinct™ MI**  
(GPU-based)



**Intel® Xeon® Series**  
(40-core)



**AMD EPYC™ Series**  
(up to 64-core)

## Embedded Domain (~10-20W)



**NVIDIA Jetson Family**  
(GPU-based)



**Kalray MPPA Coolidge**  
(80-core fabric)



**Xilinx Versal**  
(FPGA-based with DFX)

# Heterogeneous and Parallel Computing in Embedded Systems



*Performance: complex computations at high speed*



Real-time: end-to-end response time within budget



Power/Thermal: energy/temperature within budget



Safety: guarantee correctness and integrity of operation

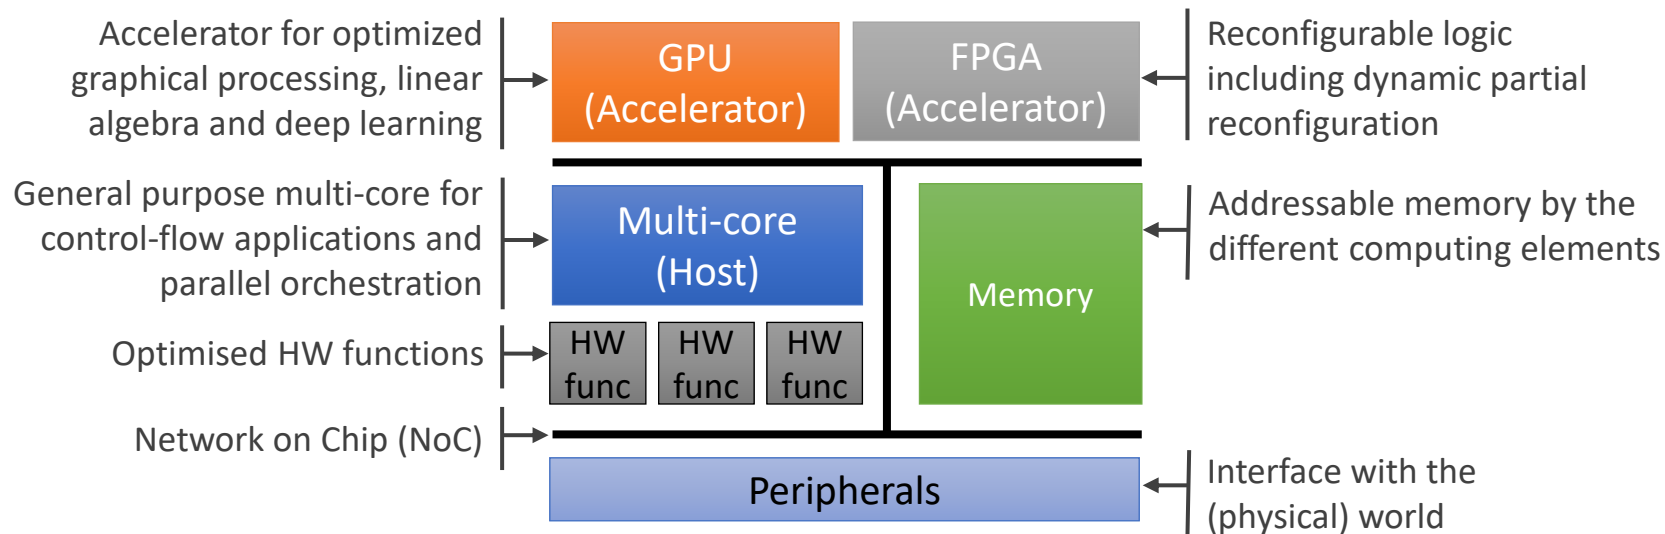


Security: prevent external elements from affecting correctness and integrity

# Heterogeneous and Parallel Computing




**Host-centric paradigm:** The parallel computation is orchestrated by the general-purpose multi-core



# The SW Productivity Gap




1. Efficiently exploit parallelism and achieve the required performance
2. Reason about the functional and non-functional correctness



Source: ITRS & Hardware-dependent Software, Ecker et al., Springer


# Addressing Complexity on System Development



## Model Driven Engineering (MDE) in Embedded Systems

1. Construction of complex systems
2. **Formal verification** of functional and non-functional requirements (NFR) with **composability** features
  - Suitable Correct-by-construction paradigm by means of code generation
3. Only for single-core execution or with **very limited parallel support**

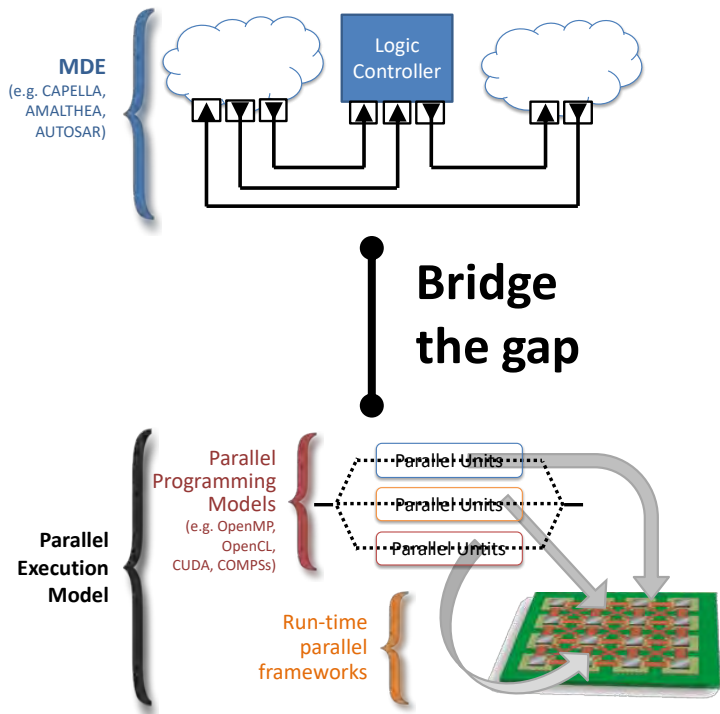
Gap between the MDE used for CPS and the PPM supported by parallel platforms



## Parallel Programming Models (PPM) in HPC

1. Mandatory for **SW productivity** in terms of
  - Programmability: Parallel abstraction while hiding HW complexities
  - Portability: Compatibility multiple HW platforms
  - Performance: Exploiting parallel capabilities of underlying HW
2. **Efficient offloading** to HW acceleration devices for an energy-efficient parallel execution

# AMPERE's Vision

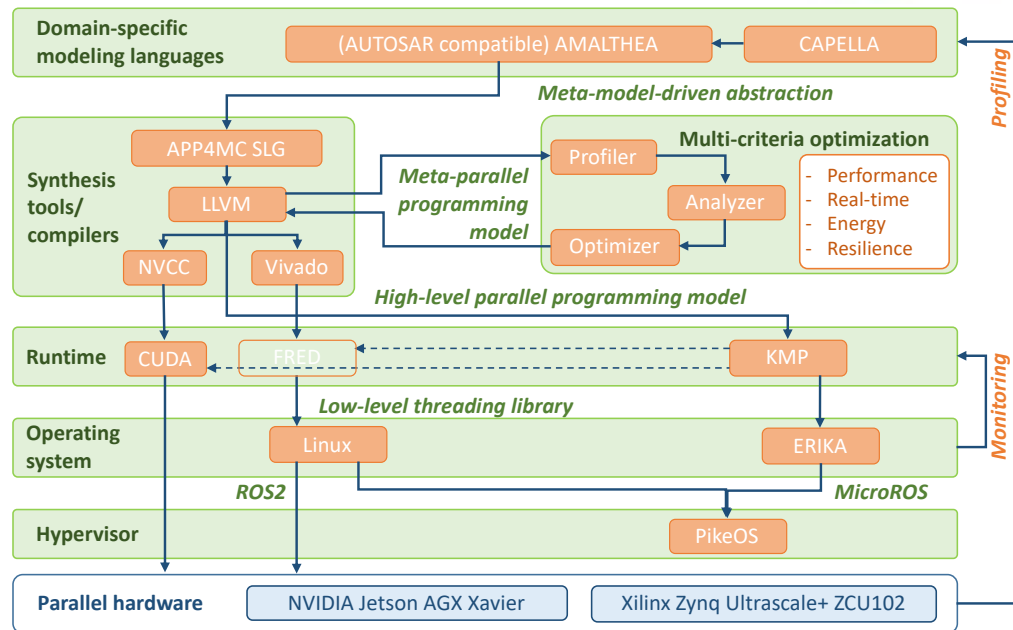


1. **Synthesis methods** for an efficient generation of parallel source code, while keeping NFR and composability guarantees
2. **Run-time parallel frameworks** that guarantee system correctness and exploit the performance capabilities of parallel architectures
3. **Integration** of parallel frameworks into MDE frameworks

# AMPERE's Main Contribution

A **novel software architecture** capable of

1. Capturing the component definition and NFR for the system model and transform it to parallel constructs
2. Fulfillment of NFR described in the CPS description
  - Real-time response, energy-efficiency, resiliency and safety and cyber-security
3. Efficient usage of advance parallel and heterogeneous embedded architectures

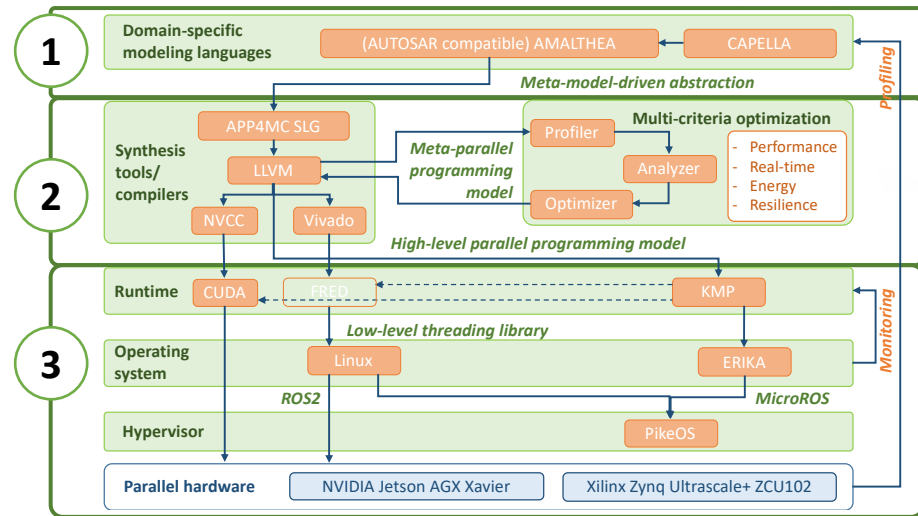
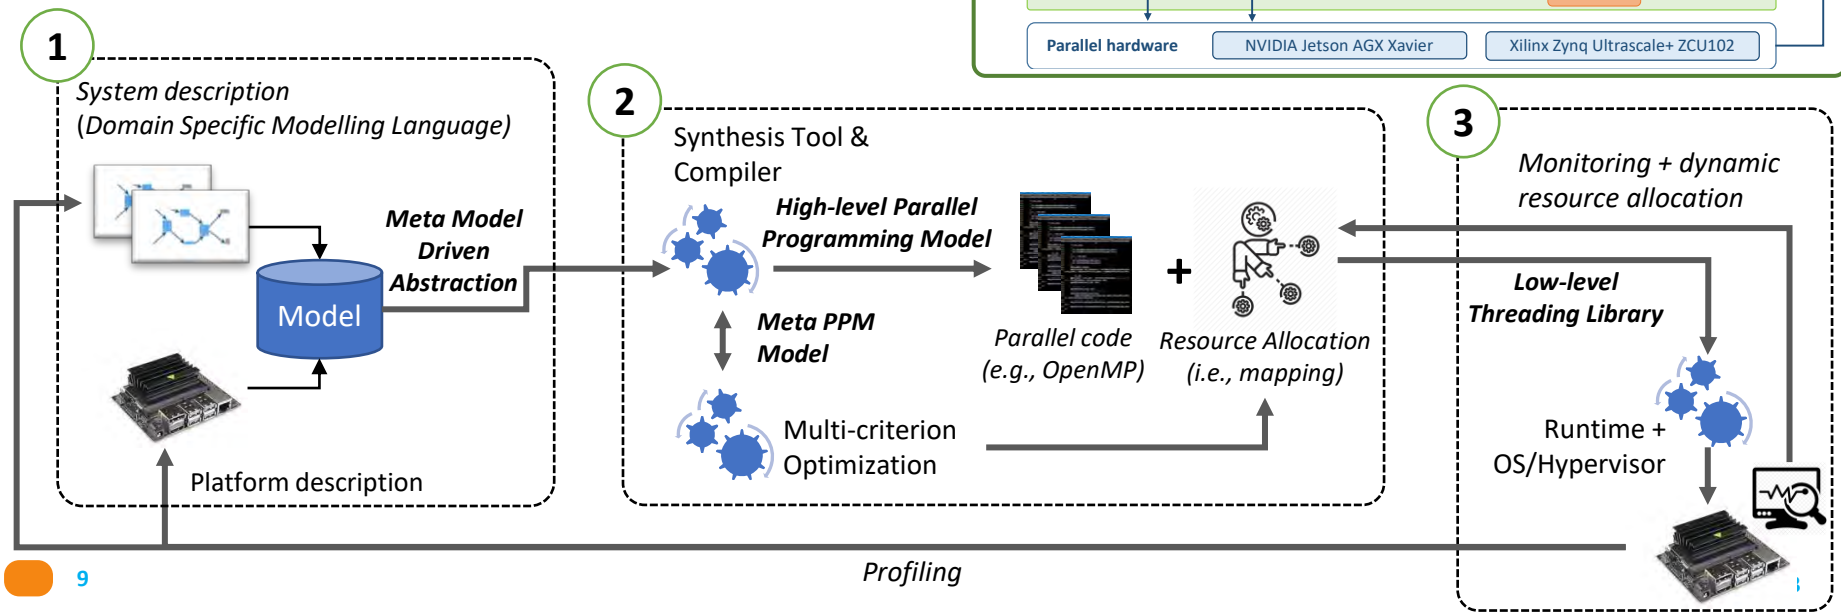


**Productivity**

- + Programmability
- + Portability/Scalability
- + (Guaranteed) Performance

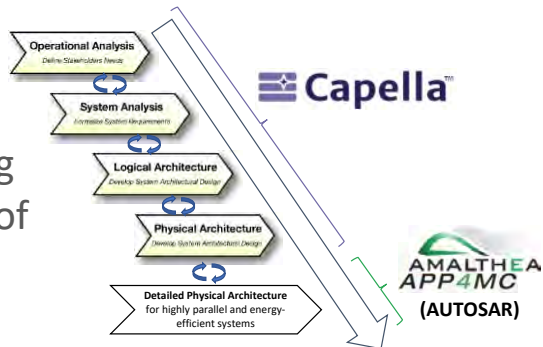


# AMPERE's Workflow Overview



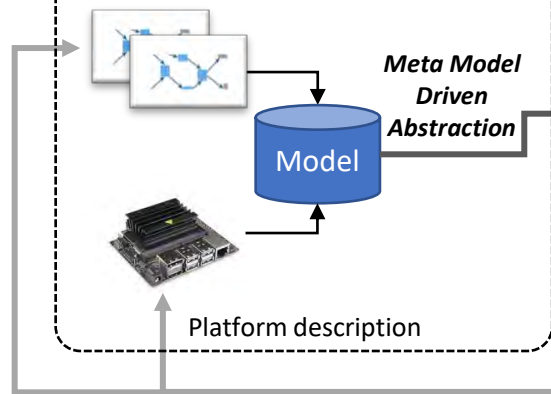
# AMPERE's Workflow Overview

DSML addressing different phases of the V-Model



1

**System description**  
(Domain Specific Modelling Language)



Synthesis Tool & Compiler

High-level Parallel Programming Model

Meta PPM abstraction

Multi-criterion Optimization

Parallel code (e.g., OpenMP)

Resource Allocation (i.e., mapping)

Monitoring + dynamic resource allocation

Low-level Threading Library

Runtime + OS/Hypervisor

Hardware Abstraction Layer

Profiling

1

Domain-specific modeling languages

(AUTOSAR compatible) AMALTHEA

CAPELLA

Meta-model-driven abstraction

Synthesis tools/compilers

APP4MC SLG

LLVM

NVCC

Vivado

Meta-parallel programming model

Profiler

Multi-criteria optimization

Analyzer

Optimizer

- Performance
- Real-time
- Energy
- Resilience

High-level parallel programming model

Runtime

CUDA

FRED

KMP

Operating system

Linux

ERIKA

ROS2

MicroROS

Hypervisor

PikeOS

Parallel hardware

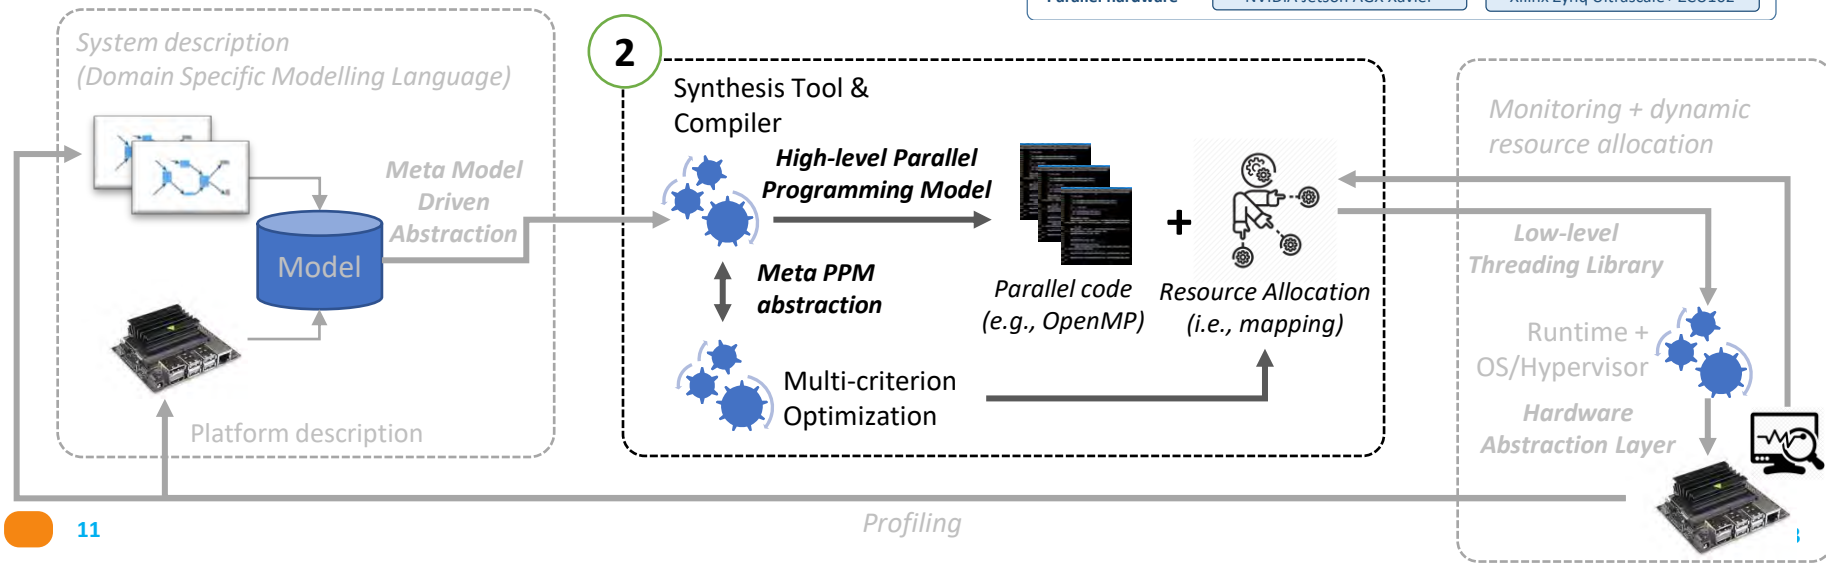
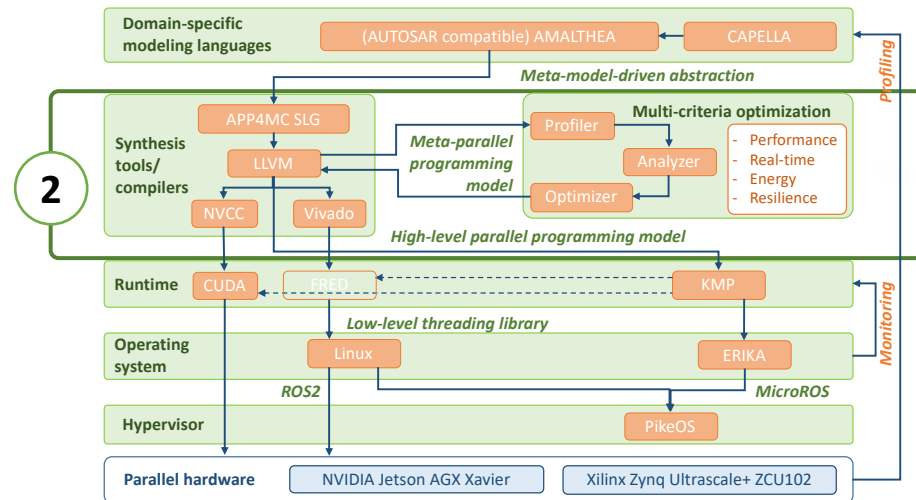
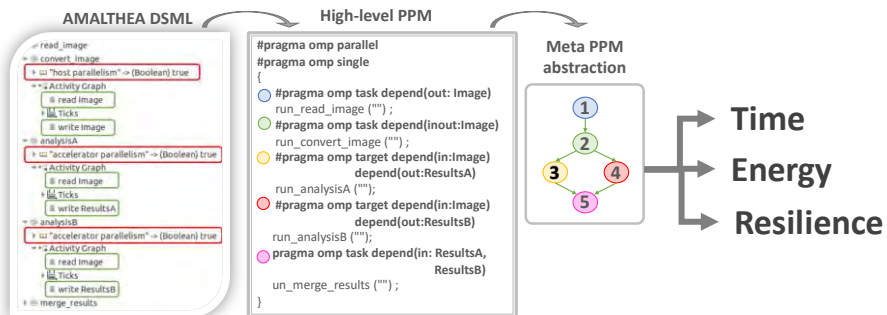
NVIDIA Jetson AGX Xavier

Xilinx Zynq Ultrascale+ ZCU102

Monitoring

Profiling

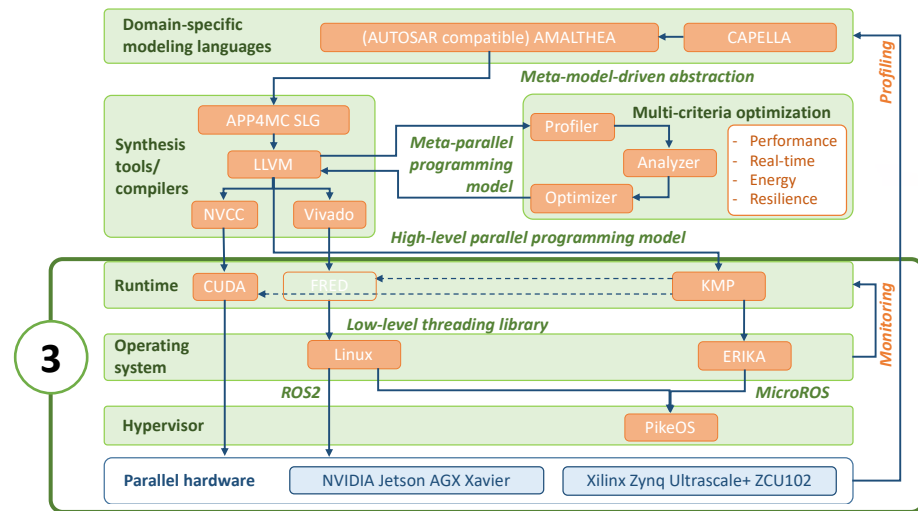
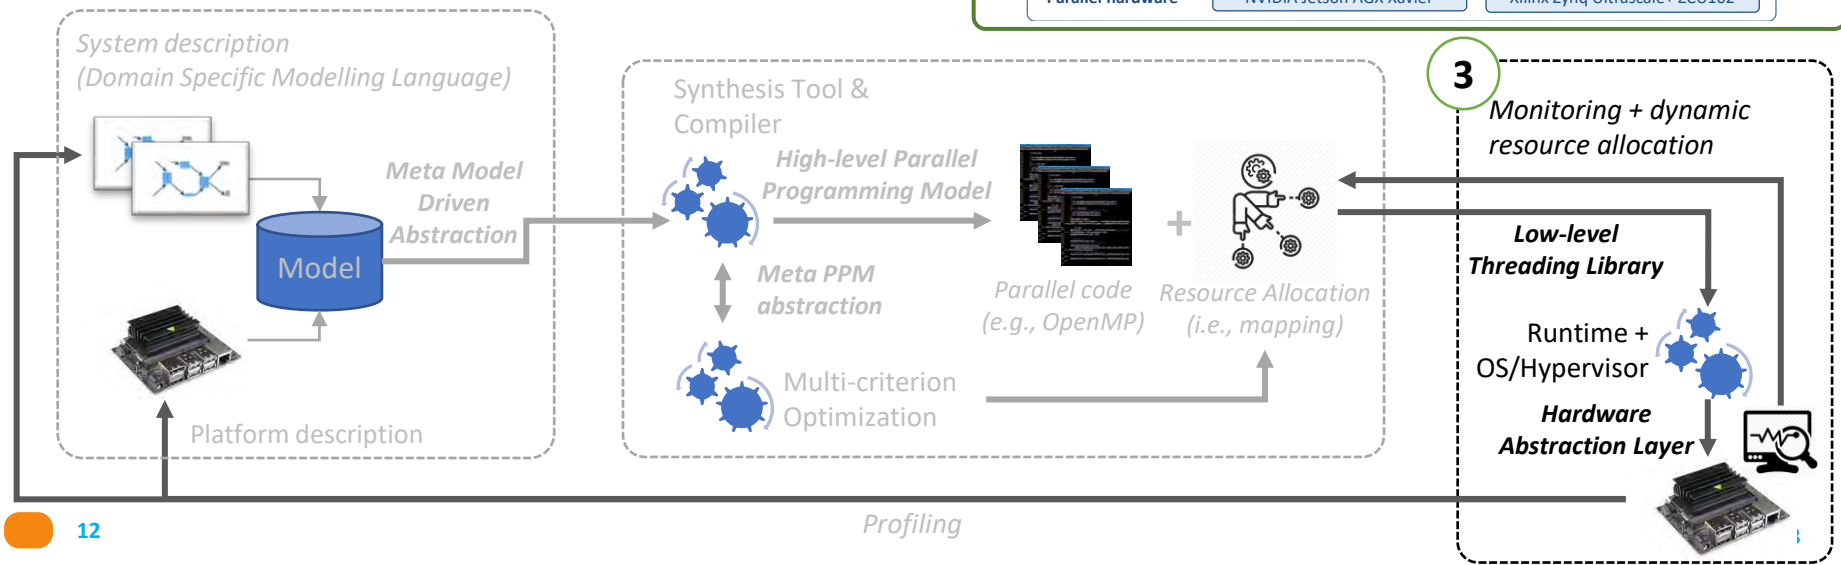
# AMPERE's Workflow Overview



# AMPERE's Workflow Overview

Runtime Monitoring  
and Implementation of

- Time
- Energy
- Resilience
- Safety and Security





## Obstacle Detection and Avoidance System (ODAS)

- ADAS functionalities based on data fusion coming from tram vehicle sensors



## Predictive Cruise Control (PCC)

- Extends Adaptive Cruise Control (ACC) functionality by calculating the vehicle's future velocity curve using the data from the *electronic horizon*
- Improve fuel efficiency (in cooperation with the powertrain control) by configuring the driving strategy based on data analytics and AI

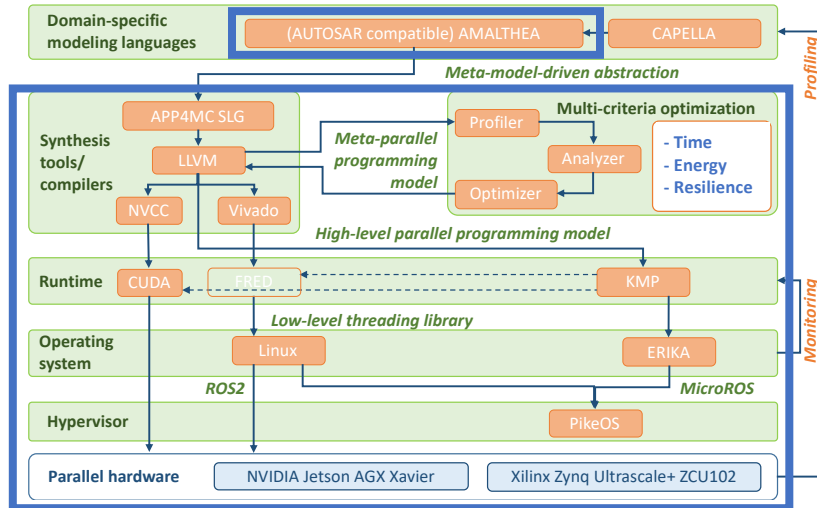


# The AMPERE Software Architecture: A Modular Design

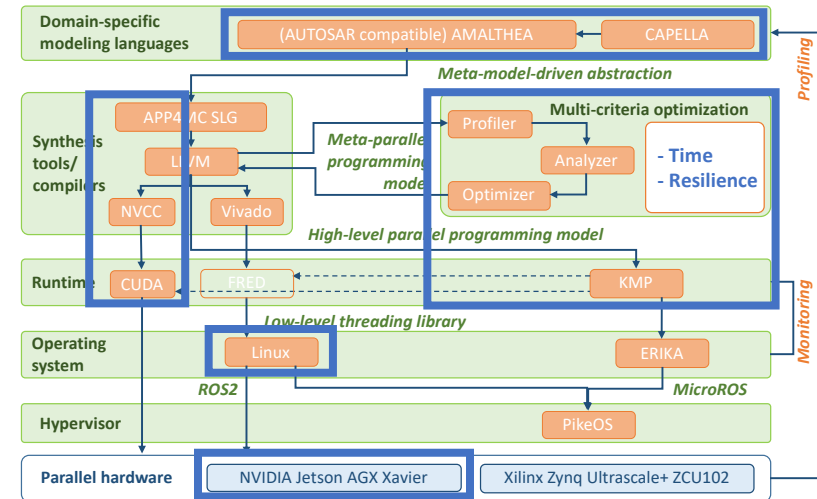


- AMPERE aims to support different “instances” of the SW architecture
  - Increase exploitation opportunities

*SW Architecture applied to PCC*



*SW Architecture applied to ODAS*



# Thank you

eduardo.quinones@bsc.es



[www.ampere-euproject.eu](http://www.ampere-euproject.eu)



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



The OpenMP logo, featuring the word "Open" in a white sans-serif font and "MP" in a larger, bold, white sans-serif font, both underlined with a white horizontal line. A small registered trademark symbol (®) is located to the right of the "MP".

# OpenMP<sup>®</sup>

## SC23 Booth Talk Series

**[openmp.org](https://openmp.org)**

OpenMP API specs, forum,  
reference guides,

and more

**[link.openmp.org/sc23](https://link.openmp.org/sc23)**

OpenMP SC23 booth talk  
videos

and presentations