Computer Architecture - Operating Systems (CAOS)

High-Integrity Systems (HIS)[*] are common in domains like transportation (automotive, avionics, railway, and space), health care, and manufacturing among others. HIS have been historically subject to several constraints in their design, implementation, validation and operation. First, **safety** both **temporal** and **functional** since any unexpected misbehavior can cause irreparable environmental damages, economic loss, harm to people and even fatalities. HIS operate under **size**, **weight**, **power**/**energy**, and cost envelops. More recently the increased connectivity and data collection performed by HIS also brings **security** concerns. Also, advanced software is increasingly used to implement control safety-critical functionality such as steering, driver assist, autonomous driving and in aerospace: flight management, mission control, and pilot displays. In fact, software is the central element for the competitiveness of HIS where implementing functionally-rich critical features is pivotal to increase product competitive edge. Hence, in the near future most advancements in these domains will be driven by sophisticated software.

? At the hardware level, heterogeneous multiprocessor systems on chip (MPSoCs) are becoming the de facto computing platform across HIS domains. MPSoCs encompass a variety of computing units like different types of CPUs and GPUs. An embedded FPGA can be used to accelerate specific functions or to provide specific support for I/O. These types of highly parallel architectures are monopolizing the market, from embedded systems to supercomputers. However, achieving high performance and compliance with functional safety standards at the software level is challenging both from a programmability point of view, as well as regarding the collection of evidence required for certification such as multicore interference, worst case execution time, absence of programming errors, etc.

? At the software level, Artificial Intelligence (AI) has already begun to show its benefits in HIS. AI can help HIS become more efficient and powerful. In fact, AI techniques like deep learning (DL), supported by the increasing computing power of modern high-performance platforms, are often the most appropriate way, if not the only way, to implement advanced software safety-related requirements for detecting objects and avoiding collisions when driving vehicles, medical diagnosis, fault detection, and security protection.

? At the modelling level, new methods are required to handle the uncertainty brought by modern and future HIS emanating from 1) the integration of complex software stacks, machine learning solutions, and high-performance computing devices; and 2) the natural sources of uncertainty that appear when the HIS interacts with the environment. Both increase the natural functional and temporal uncertainty engineers manage when developing HIS.

HIS are also referred to as real-time systems, critical (real-time systems), safety-critical systems, cyber-physical systems, and edge (AI) systems.

Objectives

In order to addressing all these challenges, the CAOS group goes beyond silo research and proposes whole-stack transformational changes, including AI algorithms (e.g. to increase the trustworthiness of AI), safety standards (e.g. to adapt to the specific characteristics of AI), hardware design (e.g. to handle hardware reliability issues and to provide guaranteed performance while adhering to safety and security requirements), software (e.g. to ease programmability and facilitate verification, validation, and testing).

Hardware

- Develop high-performance reliable, time-analyzable and low-power **processor designs**.
- Develop new **accelerators designs** for AI-based functionalities in the context of HIS leveraging constraints related to performance needs, low power operation, adherence to functional safety requirements (e.g., from ISO26262), and easing testability and certifiability.

The developed solutions range from designs demonstrated on well-established simulators to FPGA implementations, and opportunistically realizing them in ASICs (test chips), with particular emphasis on RISC-V.

System Software

- Developing **methodologies and benchmarks** to fairly evaluate different processor and accelerator designs. Understanding the strengths and the bottlenecks of current processing technologies is a key factor for their adoption in HIS industry as well as in order to make hardware proposals that are of interest to HIS industry.
- Developing **new parallel programming models** and run-time systems applicable to both computing domains, i.e. high performance computing (HPC) and real-time embedded computing (EC), to exploit the performance opportunities of the newest highly parallel embedded many-core processor architectures and accelerators such as GPUs, while providing timing guarantees and allowing software compliance with safety standards.

Modelling

- Develop **Probabilistic Analysis Techniques** for the prediction of high execution time distributions for any system at any scale (from single-cores to supercomputers), with particular emphasis on Worst-Case Execution Time for real-time HIS. The methods include fields such as Extreme Value Theory (EVT), survivability analysis, and Markov's inequality among other statistical techniques from advanced stochastic modelling.
- **AI** research on designing robust machine learning models, creating reliable AI-based systems, producing uncertainty visualization and reporting techniques to improve decision-making together with traceability. Furthermore, combine EVT techniques, which aim to enhance extreme AI forecasts and can be used in conjunction with explainability methods in order to assess the trustworthy of AI-based systems.

Validation, Verification, and Testing

- Develop and improve the **multicore microbenchmark technology** and tools to support platform characterization and timing analysis of representative heterogeneous COTS platforms in the CES industrial domains.
- Software and hardware solutions, to **ease verification and validation** of COTS HIS, as well as to implement safety measures. Those solution focus on delivering **observability and controllability** channels, as well as means to mitigate and **manage random hardware faults**, including those leading to common cause failures, in line with the safety requirements of HIS (e.g., ISO 26262 and ISO/PAS 21448 for automotive; DO178C, DO254 and CAST-32A for avionics).
- Devising software design patterns, including software specification and architecture, and their specific realizations to enable the certification of AI-based systems for safety-critical operation, such as autonomous driving or unmanned aerial vehicles.

Beyond HIS

Modelling of uncertainty using a probabilistic approach is a generic framework that also benefits other domains like financial or pharmacological ones to increase the reliability, robustness and trustworthy of any forecasting system (in particular, when that system is based on AI). By understanding the probabilistic sources of uncertainty in the forecasting system and the potential consequences of the different outcomes, it is possible to design measures to mitigate risk and ensure that system operates as intended. In addition, uncertainty modeling can help to improve the interpretability of AI-based systems by providing a framework for understanding the limitations and uncertainty of the system's output for the sake of creating Trustworthy AI-based systems.

Barcelona Supercomputing Center - Centro Nacional de Supercomputación