SAFE DYNAMIC RESOURCE ALLOCATION IN MIXED-CRITICAL SYSTEMS

Goals

Offering hypervision services to have a safe cohabitation between mixed-critical tasks handling dynamic task creation and resource allocation.

Proposed mechanisms

Minimal memory management
- separation between memory allocation and allocation policy

Dynamic mixed-critical scheduling
- bounded latency access to CPU for the most critical tasks
- asymmetric implementation of the main scheduling function → low-overhead preemptions and task switches

Access control
- small constant-time bounds on object/service access check/destruction
- no central storage → prevents DoS

Resource lending
- no memory allocation in services → dynamic scalability, DoS prevented
- resource requirements dimensioned at design time
- Bounded temporal interference even in preemtive scheduling

Prototype implementation

Bare metal kernel implementation on Kalray MPPA-256 many-core architecture

- 16 compute clusters
- 1 compute cluster = 17 Kalray cores and 2MB of SRAM
- 4 I/O clusters: 2 DDR and 1 Ethernet controllers
- Network on Chip (NoC) interconnecting clusters
- Determinism achievable under certain conditions
- Low power/performance ratio

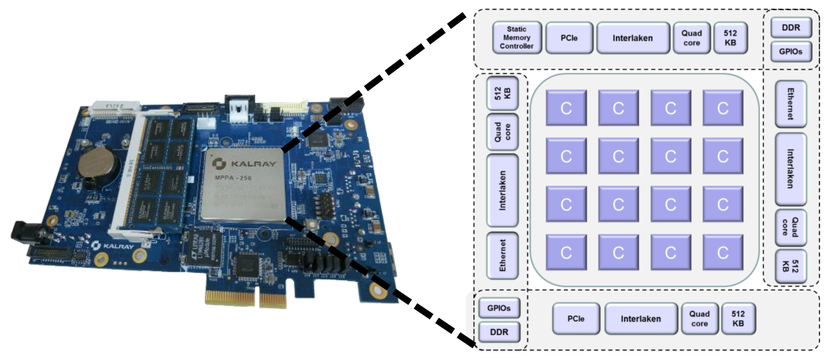

Demonstrator integration in progress

Engine Control (EC)
Hard real time constraints
DAL A

Health Monitoring (HM)
High Performance constraints
DAL E

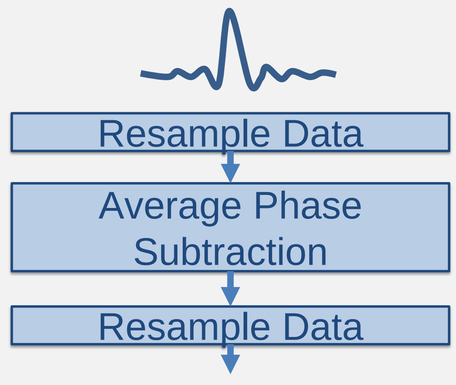

- Two applications with different criticality: DAL A instantiated statically, DAL E instantiated dynamically
- Access to shared services for I/Os (Ethernet, PCIe, NoC interfaces)
- Demonstration of steady DAL A service with and without cohabitation with resource intensive DAL E service

CONTACT | paul.dubrulle@cea.fr