AUTOMOTIVE NEEDS IN MULTI-CORE PLATFORMS

Focus on Safety and Fault-Tolerance Aspects of Mixed-IP & Mixed-Criticality Integration

Tom Fuhrman
General Motors R&D

CPSWeek 2015
RTAS Industry Panel
April 15, 2015

GENERAL MOTORS

AUTOMOTIVE INDUSTRY TRENDS

Amount of software-based functionality in vehicles is increasing

This functionality is increasingly safety-critical and security-critical (e.g., automated driving, connected vehicles)

Architectures are increasingly integrated, more and more functionality implemented on fewer and fewer nodes
(reduce size, weight, power, cost)

Software integration is increasingly mixed-IP and mixed-criticality

Hardware execution platforms moving to multi-core
(both homogeneous and heterogeneous cores)

What types of execution platforms should we use in the future? And how do we ensure desired timing / performance properties are met?

FUNCTIONALITY DRIVES THE DESIGN

Stability-enhancement systems

- Steering adjustment
- Braking adjustment

Warning-only systems

- Forward collision warning
- Lane departure warning
- Side blind zone alert
- Rear obstacle detection
- Rear cross-traffic alert

Automated driving systems

- Adaptive cruise control (throttle and braking)
- Lane centering control
- Automated lane change
- Automated parking

Crash-avoidance systems

- Emergency braking
- Emergency steering
- Emergency throttle

Mixed-criticalities are present, AND criticalities vary dynamically based on driving scenario!

ISO-26262 HAZARD ANALYSIS AND RISK ASSESSMENT DETERMINES THE CRITICALITIES

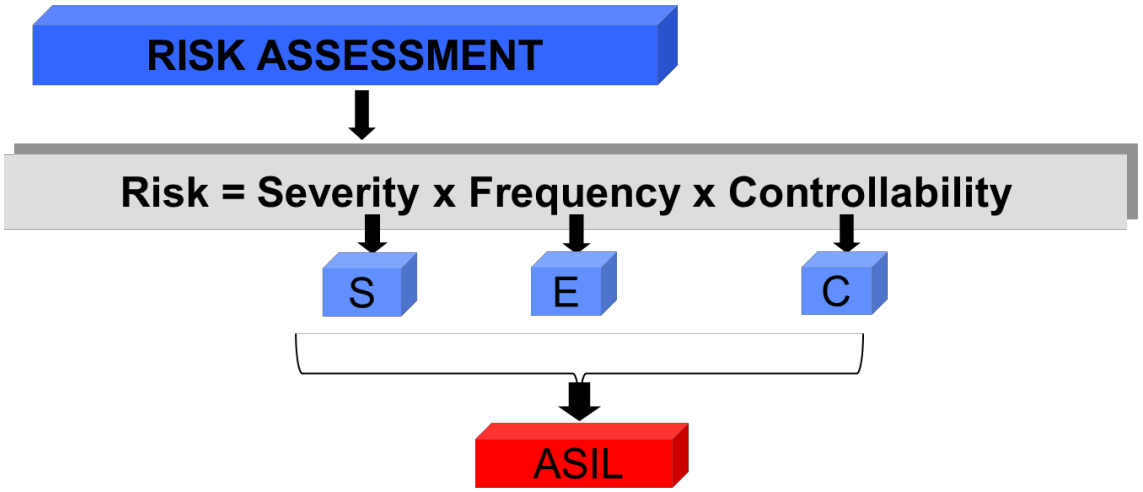

		C1	C2	C3
S1	E1	QM	QM	QM
	E2	QM	QM	QM
	E3	QM	QM	ASIL A
	E4	QM	ASIL A	ASIL B
S2	E1	QM	QM	QM
	E2	QM	QM	ASIL A
	E3	QM	ASIL A	ASIL B
	E4	ASIL A	ASIL B	ASIL C
S3	E1	QM	QM	ASIL A
	E2	QM	ASIL A	ASIL B
	E3	ASIL A	ASIL B	ASIL C
	E4	ASIL B	ASIL C	ASIL D

FUNCTIONAL SAFETY REQUIREMENTS DRIVE THE REDUNDANCY ARCHITECTURE

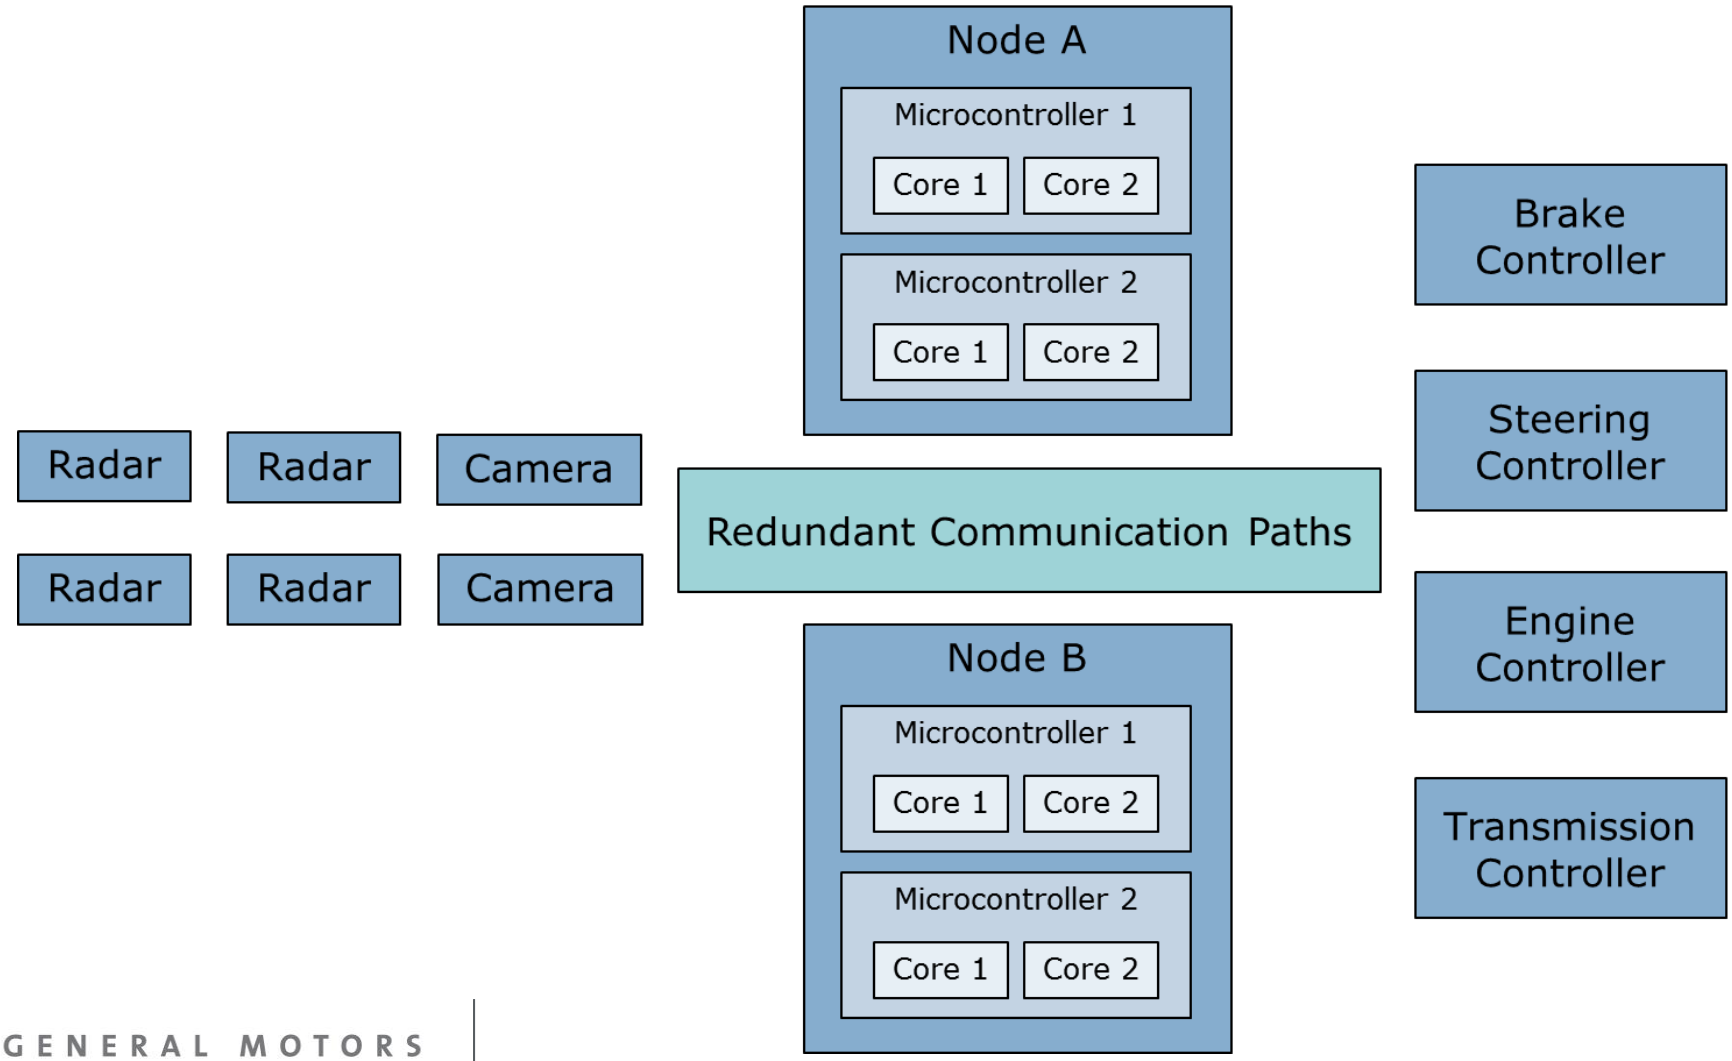

"FREEDOM-FROM-INTERFERENCE" FOR ISO-26262 MIXED-IP AND MIXED-ASIL INTEGRATION

Timing isolation techniques to consider:

- Time-slice partition scheduling (ARINC 653)
- Statically-partitioned task-to-core allocation
- Virtual machines (hypervisors)
- Dynamic task-to-core allocation with global scheduling (SMP)

BENEFITS AND CHALLENGES OF GLOBAL SCHEDULING AND DYNAMIC ALLOCATION OF TASKS TO CORES

Benefits	Challenges
• Higher processor utilization • Higher system availability under overload conditions	• Lower time-determinism • No explicit run-time fault-containment mechanism

Can timing isolation (desired timing properties of critical functions) be guaranteed by design-time analysis?

LOCKSTEP DUAL-CORE PAIRS VS. DECOUPLED PARALLEL CORES

Lockstep Dual-Core Pair

- Higher fault-detection coverage
- Lower fault-detection latency
- Comparison is easy (structural)
- Inefficient use of processing resources (all-or-nothing dual-path calculations)
- Lower power efficiency

Decoupled Parallel Cores

- Lower fault-detection coverage
- Higher fault-detection latency
- Comparison is more difficult (application-specific, or "fingerprinting" methods
- Efficient use of processing resources (selective dual-path calculations)
- Higher power efficiency

HETEROGENEOUS SYSTEM-ON-CHIP (SOC) ARCHITECTURES

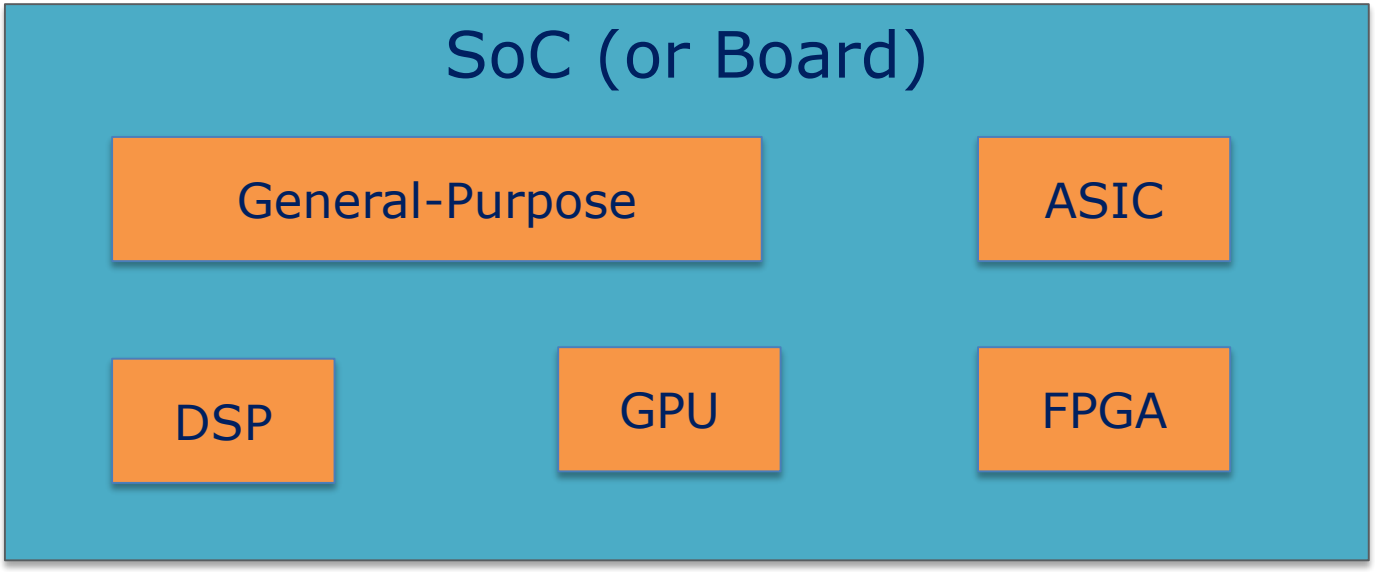

Challenges:

- System-level timing and performance analysis for design-space exploration
- Scheduling of end-to-end transactions across heterogeneous resources (must consider both computation and data flow)

DYNAMIC RECONFIGURATION ARCHITECTURES

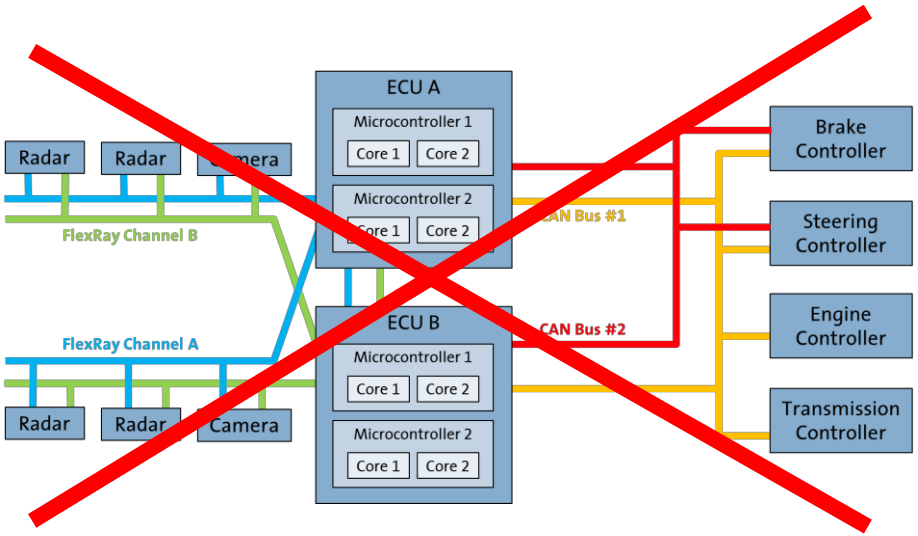

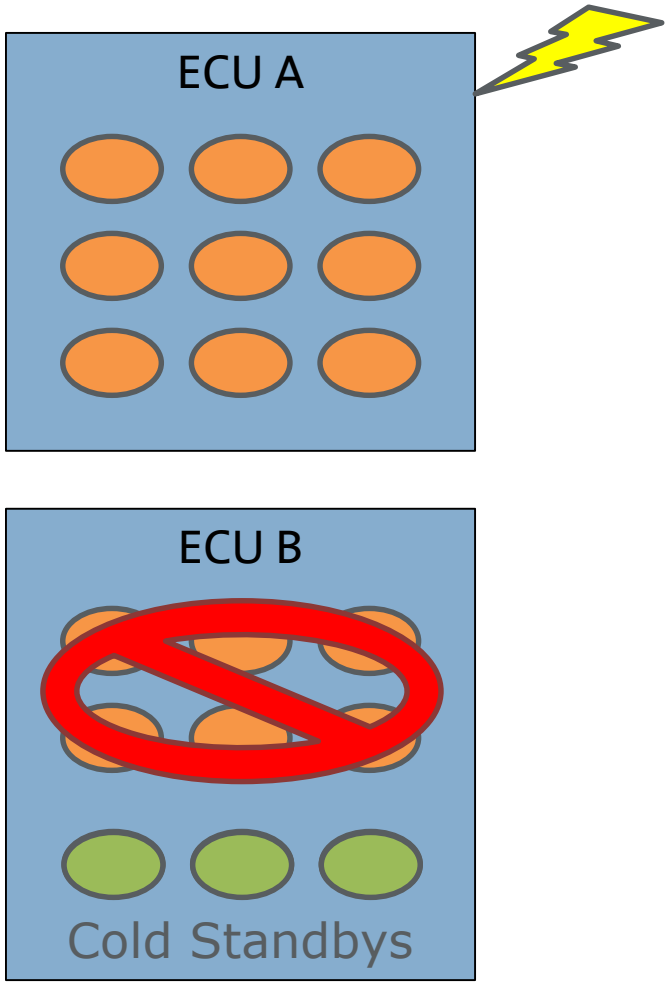

Timing and scheduling challenges:

- Cold standby activation latency
- State transfer to cold standby
- End-to-end reconfiguration timing

THANK-YOU!