FIT Analysis for Distributed Real-Time Systems

Failures-In-Time: Expected number of failures in one billion operating hours

Arpan Gujarati, Mitra Nasri & Björn B. Brandenburg

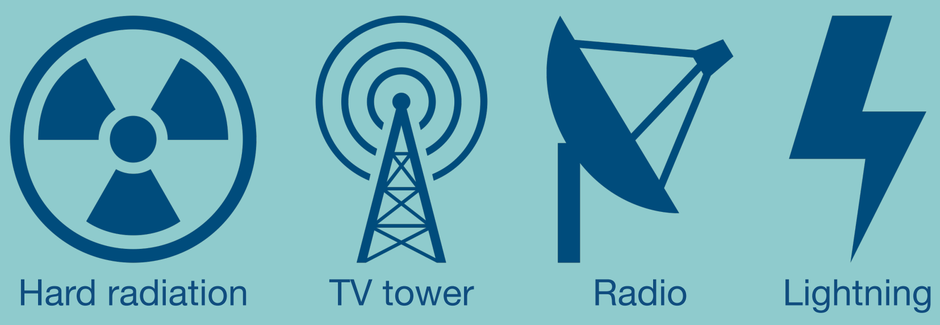

① **Transient faults (bit-flips) due to harsh environment**

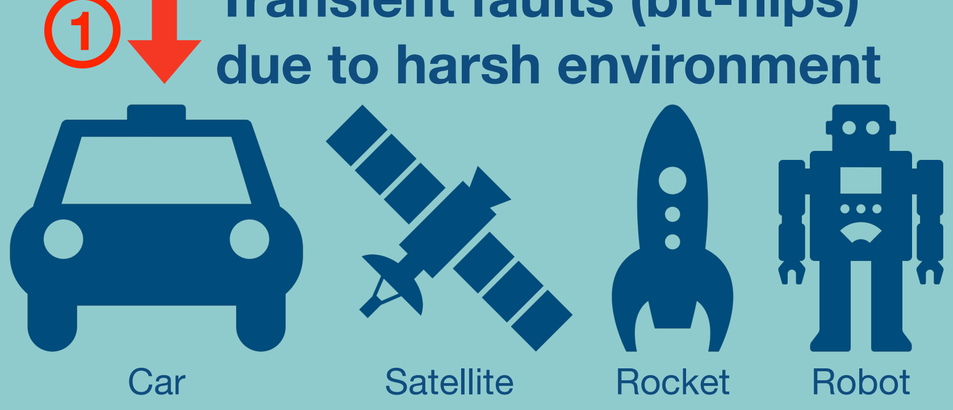

② **Time & value domain failures in time-critical embedded subsystems**

E.g., in an embedded control system:

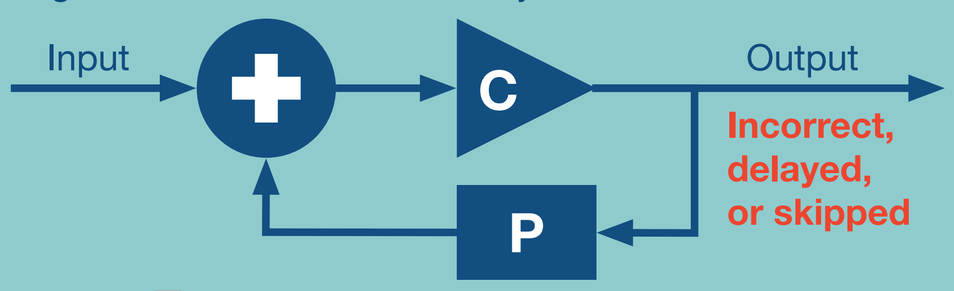

③ **Safety certification: Failure probability under a specified threshold**

As per IEC 61508 standard for electronic systems:

Zero risk of failures can never be achieved

Systems must adhere to appropriate Safety Integrity Levels (SIL), e.g.,

SIL	Continuous mode: P (failure / hour)	Low demand mode: P (failure on demand)
1	$[10^{-6}, 10^{-5})$	$[10^{-2}, 10^{-1})$
4	$[10^{-9}, 10^{-8})$	$[10^{-5}, 10^{-4})$

④ **Problem: How to quantify a safe & accurate bound on the system reliability?**

Simulation is not provably safe

Schedulability analyses only consider time domain failures

Safety and liveness proofs ignore hard timeliness

Probabilistic model checking has scalability challenges

Prior real-time analyses do not consider Byzantine errors

⑤ This Work — Provably Safe Analysis

A. All kinds of (Byzantine) failure scenarios

B. Real-time scheduling theory + Basic probability theory

C. Leverage the structure of fault-tolerant networked control systems

⑥ Model

***(m, k)-firm* model for control failure**: more than m iteration failures out of k consecutive iterations

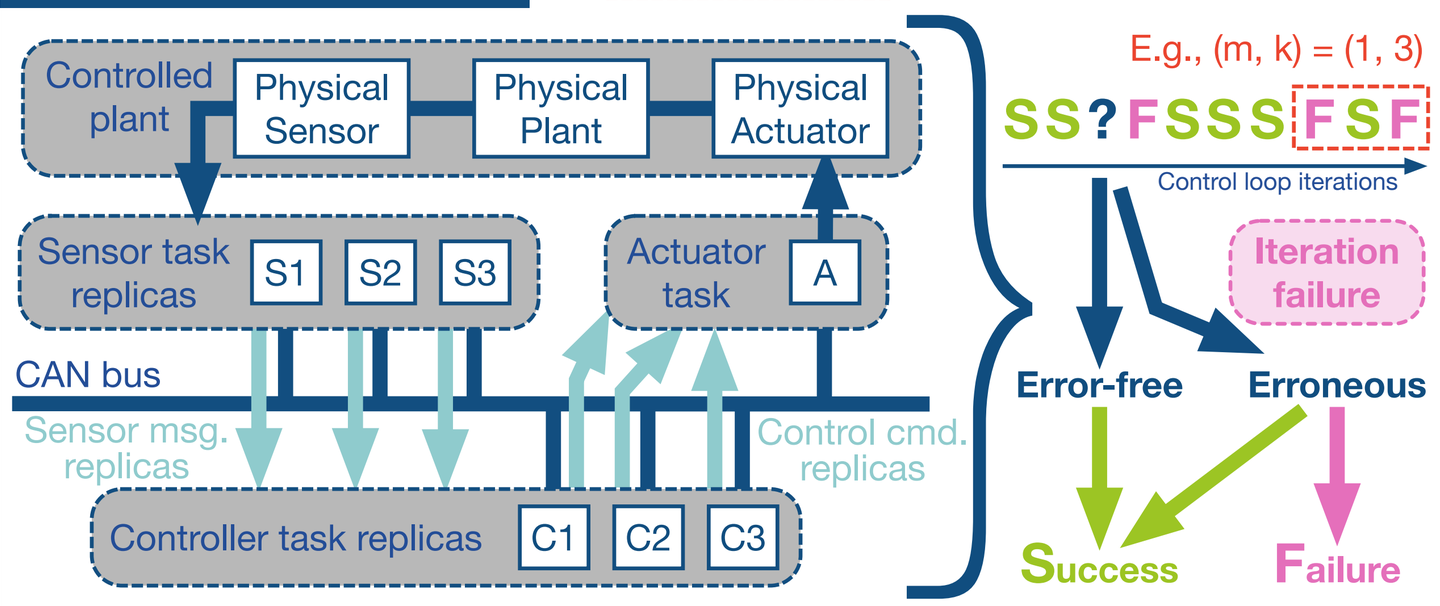

⑦ Analysis

Step 1: Upper-bound message omission, incorrect computation, & deadline violation probabilities
using peak transient fault rates derived from high interference scenarios,
and Poisson model for fault arrivals.

Step 2: Upper-bound iteration failure probability of a single control loop
accounting for interactions between different types of message errors,
and correlations due to synchronous and deterministic behavior of replicas.

Step 3: Lower-bound the Mean Time To Failure (MTTF) of the control loop
where failure denotes a violation of the plant's (m, k)-firm constraint,
using a numerical analysis technique that is both scalable and safe.

Step 4: Upper-bound the FIT rate (Failures-In-Time) for the control loop

$$\text{Upper bound on the FIT rate} = \frac{10^9}{\text{Lower bound on the MTTF (in hours)}}$$

Step 5: Upper-bound the system-wide FIT rate $\sum \begin{matrix} \text{FIT}_{\text{Loop1}} \\ \text{FIT}_{\text{Loop2}} \\ \vdots \end{matrix}$

⑧ Evaluation

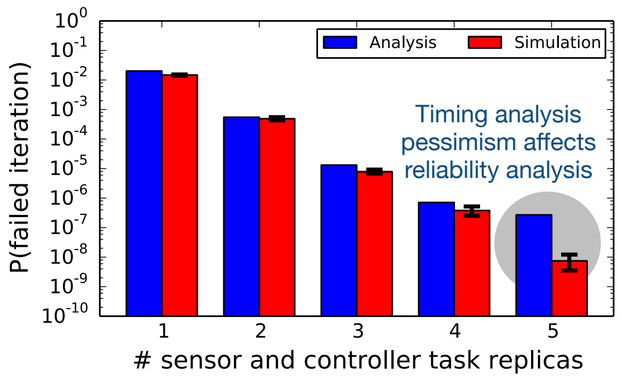

How accurate is the analysis?

When do network timing requirements (or the network schedulability analyses) become a limiting factor?

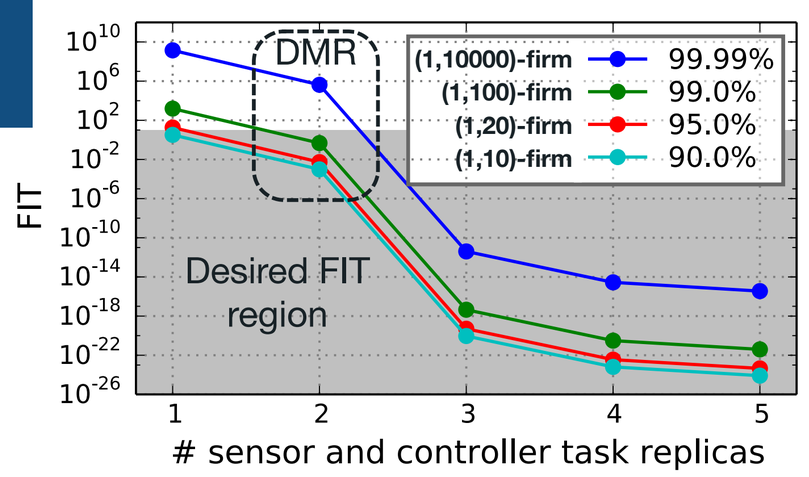

Is Dual Modular Redundancy (DMR) sufficient?

What if the desired reliability is under 1 FIT?

What if the control loop is not very robust, e.g., (1, 10000)-firm?