

# Dependable Embedded Systems — A Look Ahead

H. Kopetz  
July 2003

## Outline

- ◆ Introduction
- ◆ Hardware Developments
- ◆ Automotive Requirements
- ◆ Encapsulated Execution Environments
- ◆ Conclusion

**Dependable embedded system technology** is an important enabling technology for the industrial sector. Although comparatively small in value, this technology holds the key for determining the competitiveness of many technical products:

- ◆ Automotive (e.g., Accident-free driving)
- ◆ Aerospace (e.g., Fly by wire)
- ◆ Railways (e.g., Signalling)
- ◆ Medical (e.g., intensive care control)
- ◆ Process Control (e.g., nuclear reactors)

Because of its size, the automotive market is the most important segment for the emerging market market of dependable embedded systems.

## The 10<sup>-9</sup> Challenge in Safety Critical Applications

---

- ◆ **The system as a whole must be more reliable than any one of its components:** e.g., *System Dependability 1 FIT--Component dependability 1000 FIT (1FIT: 1 failure in 10<sup>9</sup> hours)*
- ◆ **Architecture must support fault-tolerance** to mask component failures
- ◆ Fault tolerance is based on comparing results produced within **independent fault-containment regions (FCR)**.
- ◆ System as a whole is **not testable** to the required level of dependability.
- ◆ The safety argument is based on a **combination** of experimental evidence and formal reasoning using an analytical dependability model
- ◆ **Piece to be trusted must be very small**

# Independence of Fault Containment Regions (FCR) <sup>5</sup>

---

There are two basic mechanisms that compromise the independence of FCRs in a distributed system

- ◆ Missing fault isolation
- ◆ Error propagation

**The independence of failures of different FCRs is the most critical issue in the design of an ultra-dependable system:**

- ◆ Is it justified to assume that a single silicon die contains two independent FCRs? --NO
- ◆ Can we assume that the failure modes of a single silicon die are well-behaved (e.g., fail-silent) to the required level of probability? -- NO
- ◆ How can we make sure that FCR failures are not correlated, even at a very low level of correlation (e.g., 1 in 1000)?

# 6 Independence of FCR (ii)

---


The diversity of Fault Containment Regions (FCRs) that are located on a single SoC (System on Chip) is compromised by:

- ◆ Same Physical Space (Physical Proximity Failures)
- ◆ Same Mask (Mask Alignment Issues)
- ◆ Same Bulk Material
- ◆ Same Wafer Production Process
- ◆ Same Power Supply
- ◆ Same Earthing
- ◆ Same Timing Source
- ◆

Although some of these dependencies can be eliminated, others cannot.

# Approach to Safety: The Swiss-Cheese Model

7



From Reason, J  
Managing the Risk of  
Organisational Accidents  
1997

**Independence of Layers of  
Error Detection are important**

© H. Kopetz 05.07.2003

Introduction

## A Look Back

8

In the past, many dependable embedded systems have been designed from scratch with an enormous design, development and validation effort, e.g.,

- ◆ Nuclear Control System
- ◆ Aerospace Systems
- ◆ Railway Control Systems

What is needed is an **integrated distributed architecture** and a **generic methodology for the design of dependable embedded systems supported by commercial-off-the-shelf (COTS) hardware components and software tools** such that the engineering effort needed to design, produce and validate dependable embedded systems can be drastically reduced.

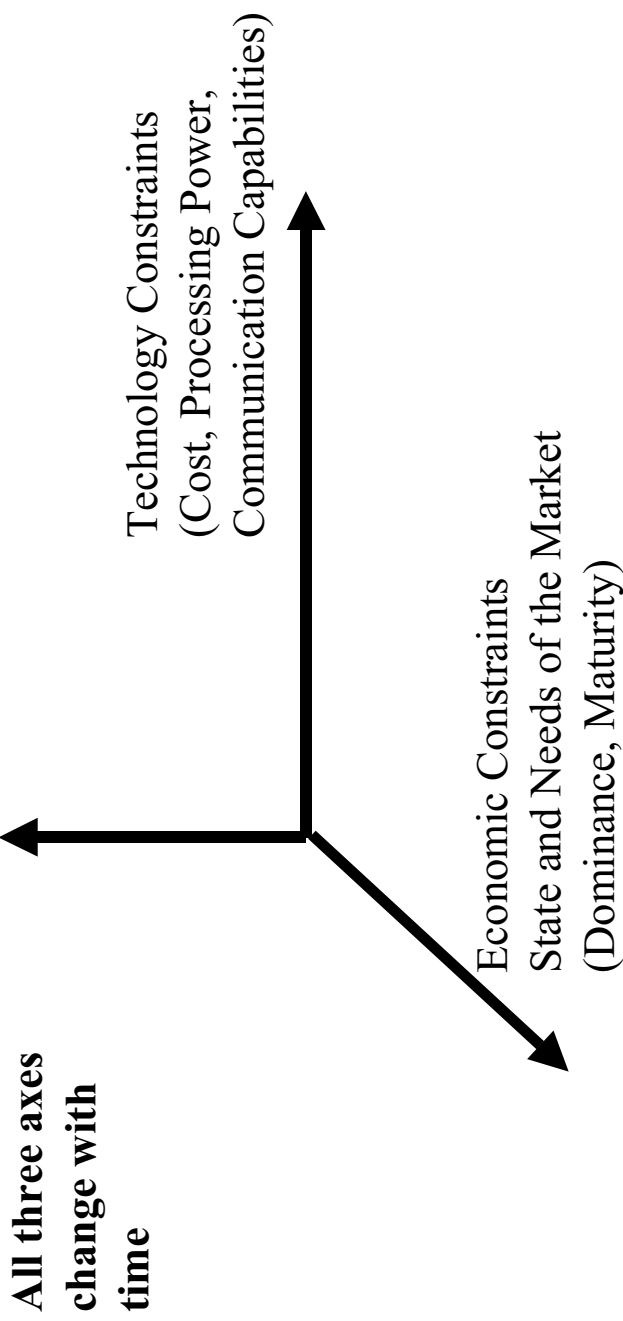
© H. Kopetz 05.07.2003

Introduction

# Window of Opportunity for COTS

---

9



© H. Kopetz 05.07.2003

Introduction

# Technology Constraints: Silicon

---

10

- ◆ At the end of this decade, we will see purely digital Systems-on-a-Chip (SOC) that will host up to one billion transistors.
- ◆ Mixed signal IC s that may include MEMS sensing and actuator elements will have a significantly lower logic density.
- ◆ From an architecture point-of-view, we will have very powerful processing nodes and smart transducers, connected via field-buses, with a limited processing power
- ◆ **In the past few years, the technological developments have accelerated. Whereas a new generation of chips is introduced every two years, it takes four years to certify a safety-critical aerospace application.**
- ◆ **Applications may live up to thirty years!**

© H. Kopetz 05.07.2003

Introduction

# Technology Constraints: Dependability

11


- ◆ The permanent failure rate of chips will be staying where it is today--between 1000 and 10 years MTBF.
- ◆ The transient failure rate will be orders of magnitude higher and is expected to increase due to reduced feature size.
- ◆ In high-dependability applications, it is not justified to assume that a single die can host more than one fault containment region: **unconstrained failure of SoCs.**
- ◆ An increasing transient failure rate (intermittent failures) are an indicator for an upcoming permanent failure.

© H. Kopetz 05.07.2003

Introduction

# Intermittent Failures: Preventive Diagnostics

12



© H. Kopetz 05.07.2003

Introduction

## Economic Constraints

---

13

- ◆ The design of a new SoC requires an investment in the order of 10 Mio € (design cost, mask costs, etc.)
- ◆ The production cost of an SoC are in the order of 10 €.
- ◆ Only applications that require millions of chips can afford the design cost.
- ◆ In the domain of dependable embedded systems only the automotive applications command a sufficiently large market.
- ◆ Europe automotive industry has a leading position in the world and thus can drive the dependable embedded systems technology.

© H. Kopetz 05.07.2003

Introduction

## Aerospace Applications

---

14

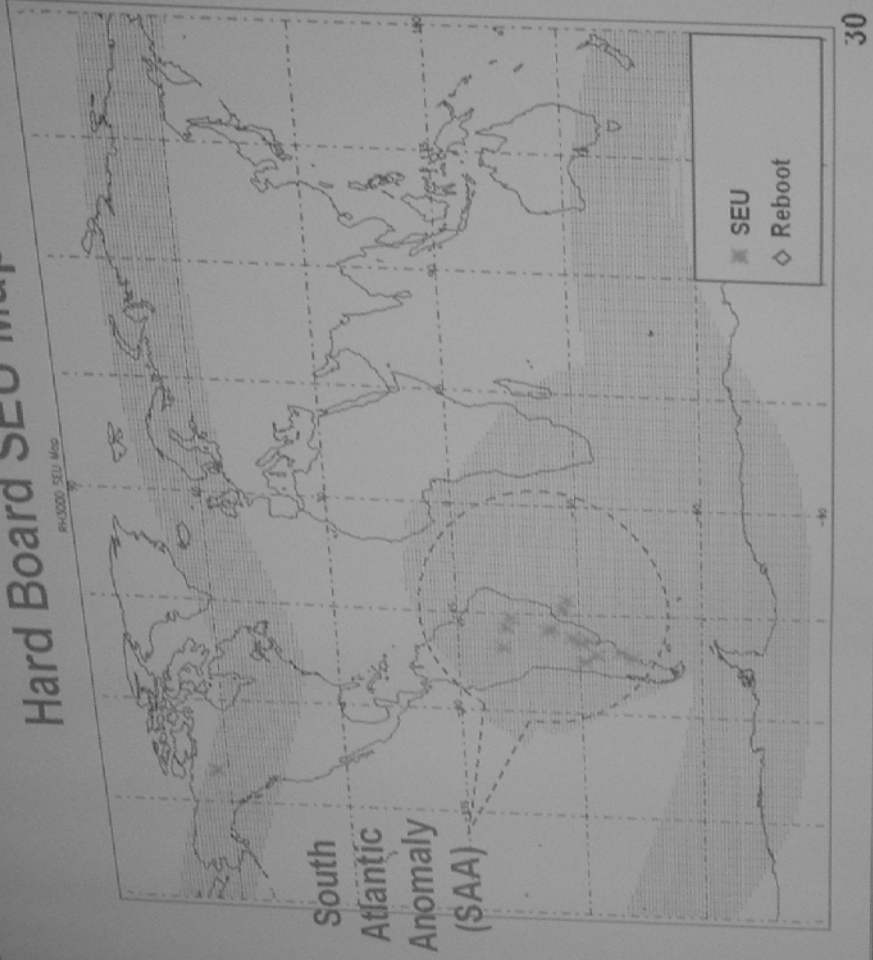
- ◆ Radiation hardened chips carry a in penalty in processing capability, power consumption and cost that is becoming difficult to justify.
- ◆ Experiments in space (e.g., the ARGOS: Advanced Research and Global Observation Satellite project) have shown that is cost-effective to use COTS chips in space and to implement the fault-tolerance by software.
- ◆ NASA is planning to use state of the art COTS components in space to perform on-board massive calculation in future scientific experiments.
- ◆ The rational for split markets in the high dependability sector is disappearing.

© H. Kopetz 05.07.2003

Introduction

1f

# Hard Board SEU Map



1c

# COTS Board SEU Map



# What can we Expect on the Hardware Side?

---

- ◆ The rate of transient failures of SoCs is on the increase due to the following
  - Single event upsets
  - Signal integrity problems
  - Variations due to manufacturing
  - Degradation problems after shipment
- ◆ The single bit-flip model is out
- ◆ Radiation hardened chips can be replaced by fault-tolerant architectures based on commodity SoCs
- ◆ The initial cost of a SoC is so high, that only applications that require millions of chips can afford their own SoC
- ◆ The pace of hardware innovations is accelerating.
- ◆ Only the automotive market is of a size that can support special SoCs for high dependability applications.

# Automotive Electronics--User Needs

---

- ◆ The wide deployment of intelligent driver-assistance systems has the potential to significantly reduce the number of accidents and to save many human lives.
- ◆ Sooner or later, *X-by-Wire* will happen. The sooner it comes, the more lives will be saved.
- ◆ The design of the *X-by-Wire* chips will be decisive, since they will constitute the *raw material* future dependable embedded systems will have to be made of.

**What are the main obstacles that hinder the wide deployment of electronic systems in cars?**

## Example of Electronics in an Upscale Car:

---

19

- ◆ Different level of controls:
  - Power train (engine, transmission)
  - Brakes, Suspension
  - Body electronics
  - Multimedia
- ◆ Federated Architecture with up to 70 nodes (Electronic Control Units--ECUs) in an upscale car
  - Essentially, every new function requires a new box
- ◆ Different networks
  - LIN fieldbus (< 20 kbits/s)
  - CAN (< 500 kbits/s)
  - MOST (Multimedia > 10 Mbits/s)

## What is Different in the Automotive Industry?

---

20

- ◆ Large number of cars (50 million/year)
- ◆ Minimization of recurring costs in a mass market
- ◆ Very high level of dependability at affordable cost
  - Majority of recalls are hardware related failures
- ◆ Few independent automotive companies in the world
  - Large enough to make their own COTS
- ◆ Attitude: *We own the world* --and in some respects they do
  - Example CAN
  - Convergence Conference on Automotive Electronics
  - Absence of academics at relevant SAE meetings (e.g. Naming)
- ◆ Difficulties when it comes to interfacing with the worldwide information infrastructure: example MOST

Why don't we move ahead?

After discussions with automotive companies, we have identified the following five major current obstacles:

1. Electronic Hardware Cost
2. Diagnosis and Maintenance
3. Dependability
4. Development Cost: Limited Reuse
5. Intellectual Property (IP) Protection

## Electronic Hardware Cost

---

Hardware costs are recurring costs that are decisive for the economic success in a mass market.

- ◆ At present, the electronic architecture on-board vehicles is *federated*, not *integrated*.
- ◆ In a federated architecture every new function requires a new electronic box (ECU-Electronic Control Unit).
- ◆ Today we find more than 70 ECUs in upscale cars.
- ◆ In an *integrated* architecture the number of hardware boxes can be reduced significantly, resulting in a significant reduction of the hardware costs.
- ◆ **The technology to support an integrated architecture with encapsulated execution and communication services is not yet mature.**

## Diagnosis and Maintenance

---

23

- ◆ The vast majority of failures in the electronic system of a car is *transient* or *intermittent*, but not permanent.
- ◆ The present electronic architectures within cars do not support the diagnosis of transient faults in an optimal way.
- ◆ The ratio of *first-time-correct* maintenance actions is in many scenarios below 50 %.
- ◆ If we assume that 2% of the cost of a car (300 € per car) are spent for electronic diagnosis, the world-wide automotive electronic diagnosis market is 15 000 000 000 €.
- ◆ **The technology to diagnose correctly transient malfunctions needs to be developed further.**

© H. Kopetz 05.07.2003

Introduction

## Dependability

---

24

- ◆ According to the ADAC statistics in Germany close to 50 % of the failures of cars on the road are caused by defects in the electronic systems.
- ◆ Connector failures are an important failure class.
- ◆ Fail-operational applications (e.g., X-by-Wire) require a reliability that must be better than the reliability of the mechanical system they replace--a level of electronic system safety that the automotive industry is not yet used to.
- ◆ The aircraft industry has the longest experience in designing safety-critical *by-wire-systems*.

© H. Kopetz 05.07.2003

Introduction

## Development Cost

---

- ◆ The unintended side effects between different application subsystems increase significantly the development and integration efforts.
- ◆ There is only a limited reuse of software and existing IP due to the missing composability support of current electronic architectures.
- ◆ The hardware environment changes so quickly, that it is difficult to consolidate the application development.
- ◆ **As a consequence, modular development, validation and certification are still more on the wish-list than in the real world.**

## Intellectual Property (IP) Protection

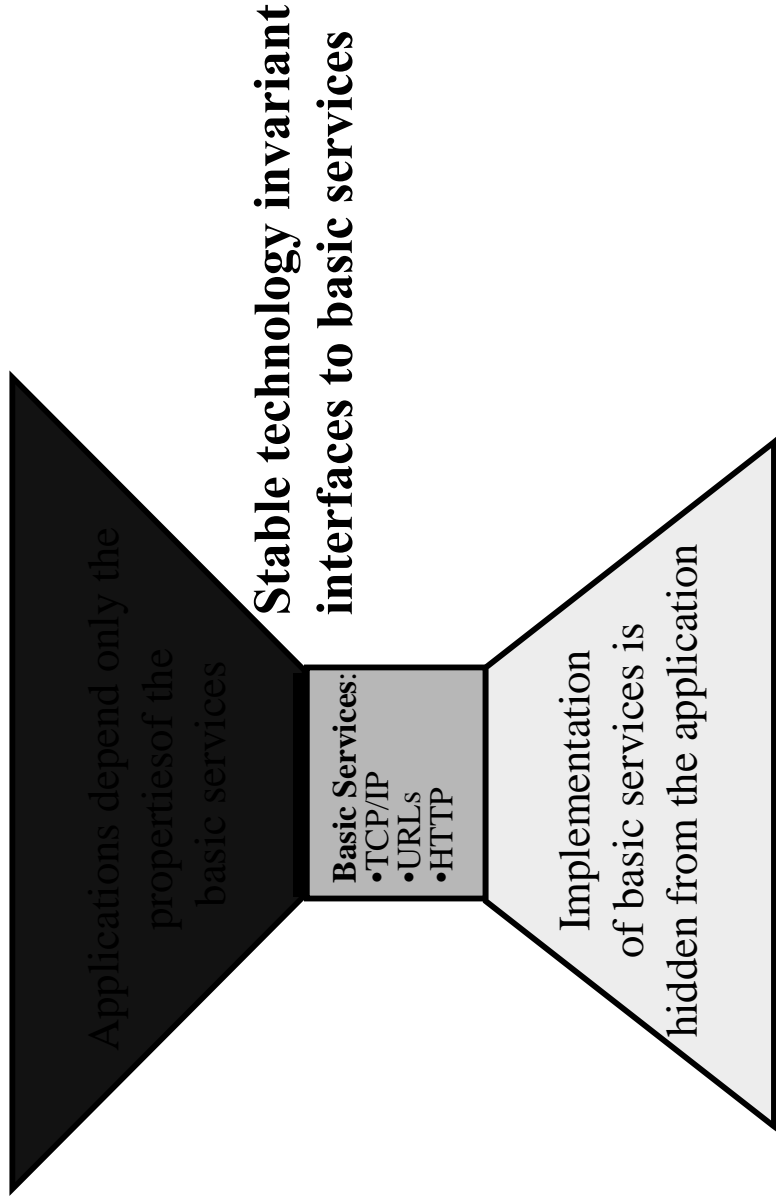
---

- ◆ Sub-suppliers of the car companies are not very willing to open their IP, because they are afraid of giving up their competitive edge (e.g., software for engine control).
- ◆ Without a deep knowledge of the software-internals, car companies are reluctant to accept system responsibility for the correct operation of ECUs that contain software modules from different sub-suppliers.
- ◆ **The contractual and legal implication of fault-diagnosis and repair responsibility of multi-vendor ECUs are difficult to resolve.**

## What makes Internet such a Success?

---

27




## Needed: An Integrated Distributed Architecture

---

28

An **Integrated Distributed Architecture** for dependable embedded applications is badly needed:


- ◆ In an Integrated Distributed Architecture the number of nodes (ECUs) can be significantly reduced by providing **multiple encapsulated execution environments** for different *Distributed Application Subsystems* (DAS) that are integrated within a single physical node and protected from each other.
- ◆ The number of cables and connectors can be reduced by providing **multiple encapsulated virtual networks** on a single wire.
- ◆ **Generic services for strong fault isolation, fault tolerance and monitoring** are provided at the architecture level.
- ◆ **Standard technology invariant interfaces** are provided by the middleware to the application, irrespective of the physical current physical environment, which can evolve.



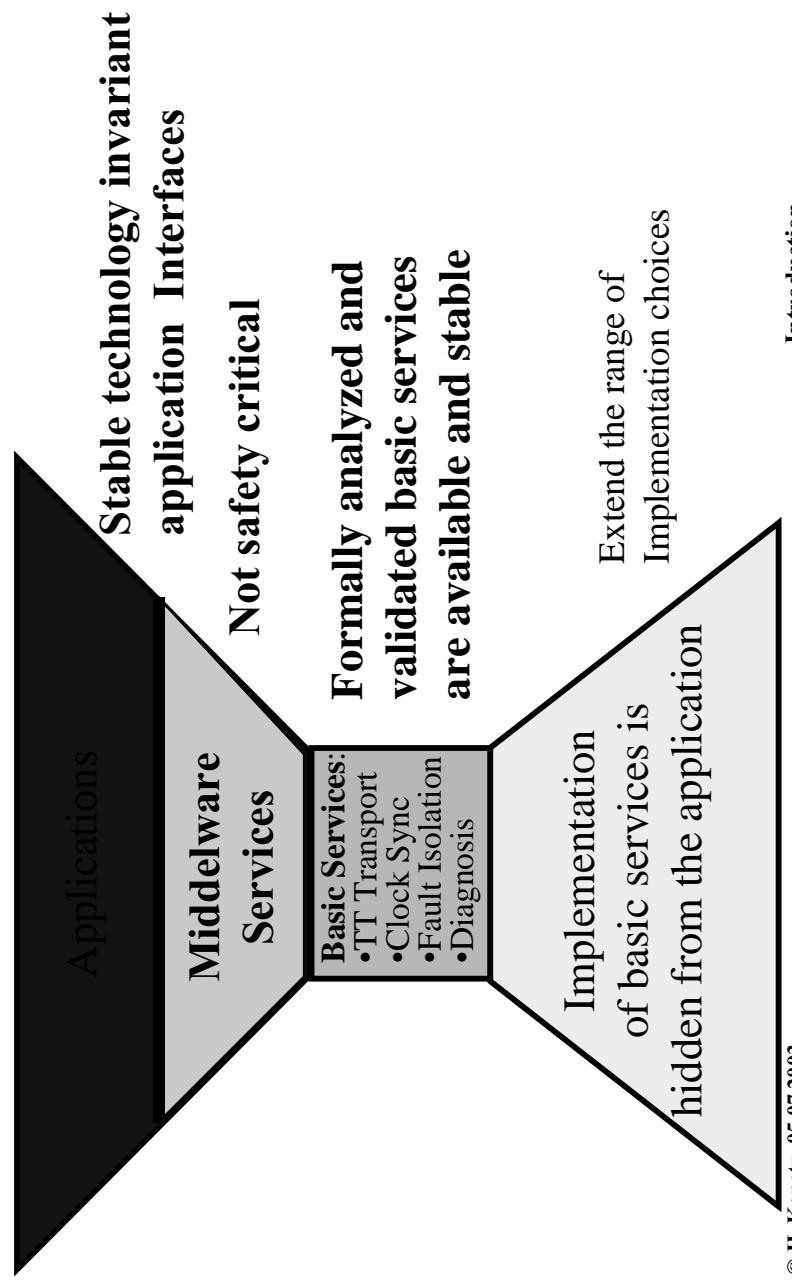
## Multiple Encapsulated Execution Environments

- ◆ Are encapsulated by middleware and **do not interfere** with each-other, neither in the domains of time no value.
- ◆ Provide **standardized technology invariant interfaces** to the multiple distributed application systems (DAS).
- ◆ The specified operation of a DAS is **continuously monitored** by an architecture based diagnostic service.
- ◆ Supports the **free movement of applications** within a single DAS (load sharing within the nodes).
- ◆ **Provide strong fault-isolation of the applications** determined by **basic services** of an architecture.
- ◆ Fault tolerance is provided by the **services of the middleware.**

# Architecture Vision for Embedded Systems



# Technology Invariant Interfaces



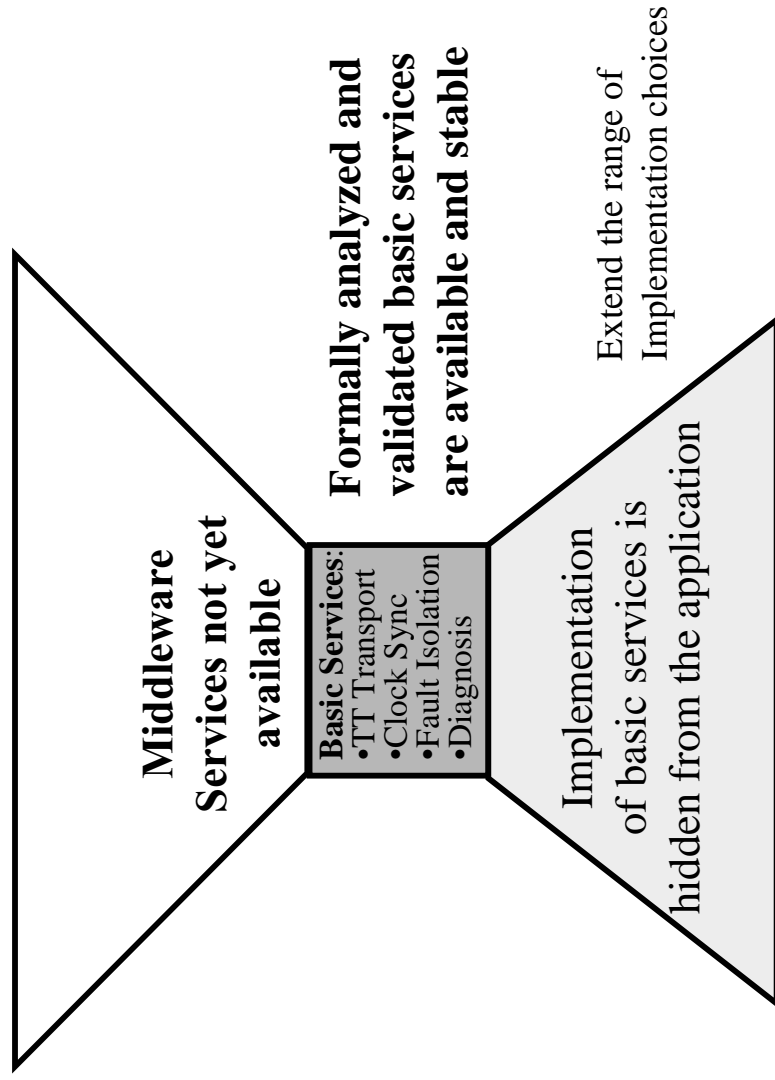
**Critical and Non-critical services can only be integrated within a single architecture if the architecture supports the safety requirements of the most critical service class.**

They must be separated for the following reasons:

- ◆ The basic services **guarantee fault-isolation and independence of FCRs.**
- ◆ The **basic services must be minimal** in order that their correctness can be established.
- ◆ The middleware services are not in the same criticality class as the basic services, since it must be **assumed that an SoC can fail in an arbitrary failure mode.**

## The TTA Provides the Basic Services

---



TTP/C-1-based hardware prototype with XILINX 600k FPGA developed within NEXT TTA(tested by IST project FIT):

## Heavy Ion Experiments (at Chalmers):

Bus topology: 37036 faults--78 error propagations (0.21 %)

Star topology: 26600 faults-- 0 error propagation

## Software Implemented Fault Injection (Vienna):

Bus topology: 562122 faults--14 error propagations (0.02 %)

Star topology: 541744 faults-- 0 error propagation

Published at DSN, San Francisco, June 2003

Formal Verification using Model Checking (SAL, UPPAAL2k) and Theorem Proving (PVS) is ongoing in the NEXT TTA Project.

## Conclusion

---

- ◆ Hardware gets **more powerful but less reliable** at an accelerating pace. The reasons for split hardware markets are disappearing.
- ◆ We must move from the **federated architectures** of today towards the **integrated architectures** of tomorrow that provide **stable technology invariant interfaces** to the applications.
- ◆ The full integration effect can only be achieved, if the base architecture supports the **highest criticality class**.
- ◆ The **Automotive Market** is the driver of the market for dependable embedded systems.