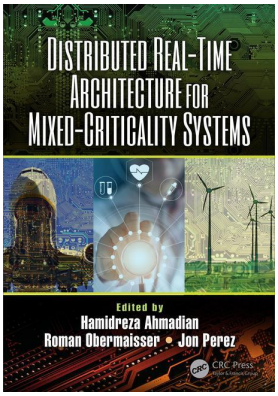

Distributed Real-Time Architecture for Mixed-Criticality Systems

Hamidreza Ahmadian, Roman Obermaisser, Jon Perez

State-of-the-Art and Challenges

Publication details
https://www.routledgehandbooks.com/doi/10.1201/9781351117821-3
H. Ahmadian, M. Coppola, M. Faugére, D. Gracia Pérez, M. Grammatikakis, I. Martinez
Published online on: 21 Aug 2018

How to cite :- H. Ahmadian, M. Coppola, M. Faugére, D. Gracia Pérez, M. Grammatikakis, I. Martinez. 21 Aug 2018, *State-of-the-Art and Challenges from:* Distributed Real-Time Architecture for Mixed-Criticality Systems CRC Press
Accessed on: 14 Jun 2021
https://www.routledgehandbooks.com/doi/10.1201/9781351117821-3

3

State-of-the-Art and Challenges

H. Ahmadian
Universität Siegen

D. Gracia Pérez
THALES Research & Technology

M. Coppola
STMicroelectronics

M. Grammatikakis
TEI of Crete

M. Faugére
THALES Research & Technology

I. Martinez
IK4-Ikerlan

This chapter provides an overview about the state-of-the-art solutions in three mixed-criticality industrial use-cases. It discusses the research challenges of three use-cases with respect to system architectures, platform technologies and development methods. Section 3.1 is dedicated to the avionics domain and introduces the modular architecture as the state-of-the-art solution in this domain. The challenges are introduced, as the performance and safety requirements of the domain are increasing. It is described how the demanded performance and support for mixed-criticality domains are achieved by the concept of partitioning and at the end the concept of security is discussed. Wind-power is the next industrial domain that is described in Section 3.2. In this section, the state-of-the-art solutions for control and supervision systems of wind-turbines are elaborated. Integration and flexibility, support for mixed-criticality, safe communication and safety certification are identified as the main challenges for such systems. At the end, the challenges of the health-care domain are elaborated in Section 3.3. Wearable intelligent devices are identified as the state-of-the-art solutions in the smart hospitals. In such solutions, security and functionality are identified as the key research and technological challenges.

3.1 Avionics Domain

The avionics domain is one of the industrial domains, in which safety has been the main driver for the development of new solutions, particularly in the civil domain. Both governmental and industrial bodies have defined regulations and standards to develop the exploitation of avionics systems. Regulations are legal documents and are the basis to define the certification of the system, if the system can operate or not. Regulation bodies are typically state agencies, like the Federal Aviation Administration (FAA) in the United States or the European Aviation Safety Agency (EASA) in the European Union. Standards are typically industry practices accepted by regulation bodies as means to develop products that follow the regulations, and in some cases as solutions to ensure the interoperability of products components.

From the technical perspective, two standards are used as basis for the development of aircraft computation systems: DO-178C [3] and RTCA DO-254/EUROCAE ED-80 [61]. These standards define the Development Assurance Level (DAL) and categorize the effects of a failure condition in the safety of the aircraft. Based on the DAL, these standards define the properties and development process that a system must follow. Table 3.1 illustrates the defined DALs by the standards, their associated failure condition in case of occurrence and the maximum accepted occurrence rate.

Table 3.1: Development Assurance Level (DAL), the Associated Failure Condition in Case of Occurrence and the Maximum Accepted Occurrence Rate

Level	Failure condition	Failure rate
A	Catastrophic	10^{-9}/hour
B	Hazardous	10^{-7}/hour
C	Major	10^{-5}/hour
D	Minor	10^{-3}/hour
E	No effect	n/a

The RTCA DO-178/EUROCAE ED-12 standard defines two properties that need to be ensured in a safety-critical system (i.e., with DALs from A to D): spatial partitioning and temporal partitioning. Spatial partitioning refers to the capability of a system to ensure that an application data is isolated in the system, i.e., it cannot be read or modified by external entities and the application cannot access other data than its own. In modern systems, temporal partitioning is often established by configuring the processor Memory Management Unit (MMU) to ensure the spatial partitioning requirements, typically by the operating system software.

Temporal partitioning refers to the capability of a system to allocate time windows to an application, during which the application is executed without being interrupted, and the application will not be executed outside the allotted time windows. Typically, temporal partitioning is established by the operating system or a hypervisor by implementing cyclic scheduling with time windows allocation in the scheduling cycle.

Avionic products must comply with the regulation of the legal bodies of the countries, where they will be used. Following the above standards during the development process allows the product developers to ensure that their products can be certified, commercialized and used. It must be noted that standards provide sufficient guidelines to ensure that the development of a product complies with the regulations, however, product developers might use other approaches than those proposed in the standards and still certify their products. The certification of products that are developed using practices other than those proposed by the standards are typically longer to certify, as the developer has to provide additional proofs that the product complies with the regulations, while in case of using the standard guidelines, the proofs are already provided by the standard.

3.1.1 State-of-the-Art: Integrated Modular Avionics

Similar to automotive domain, the electronic solutions have been introduced in the avionics domain to improve the safety of the aircraft and to reduce their weight by removing mechanical components. Initially, each functionality (or application) was typically delivered by a different single computing unit (or Local Resource Scheduler (LRS)). As a consequence, in case communications between different functions were required, safe networks, like AFDX, were deployed. This approach is known as *Federated Avionics Architecture* or *Federated Avionics System* (*Federated Architecture* or *Federated System* for short).

While the federated design has been successfully used for the development of several generations of aircraft products, the increase in usage of computing solutions for implementing the avionics functions has become a challenge in terms of *Size, Weight, Power and Cost (SWAP-C)*.

Integrated Modular Avionics (IMA) [62] was introduced as a solution to combine functionalities in a computing unit and to reduce the number of computing units in an aircraft. Basically, this solution involves an enhanced operating system or hypervisor, which ensures disjoint time windows (defined by the system integrator) in the cyclic schedule for each of the functionalities combined in the system. Furthermore certain IMA approaches enable the application of incremental certification: for a certified IMA system that combines functions A and B, the addition of a new function (e.g., C) does not require a new certification of the entire system, but only the added function. The IMA is standardized in the RTCA DO-297/EUROCAE ED-124 standard [62] and also considers the modularity of the hardware parts of an embedded system.

3.1.2 Challenges

The demanded performance and safety requirements in avionics are continuously increasing. The SESAR [63] project has the objective to to increase its capacity of the European air traffic by modernizing it, through the development and implementation of new procedures and technologies impact both ground control and the aircraft themselves. The Clean Sky [64] project develops technologies to reduce CO^2, gas emissions and noise levels produced by aircrafts. The solutions proposed by those programs require embedded processing systems with increased performance to support their implementations. Furthermore, the ongoing trend towards the usage of avionics solutions in urbanized areas, like UBER with their on-demand urban air transportation project Elevate [65], the Dubai flying taxi trials [66, 67] or the Airbus Vahana project [68] to provide urban air mobility solutions, introduces new and major challenges. To target this new market, embedded systems do not only need to provide more performance, but also more integration, i.e., able to combine multiple applications in a single device.

With the new performance and safety requirements introduced by these projects, the expectation on the embedded systems capabilities is increased, forcing the industry to move away from the single-core solutions. Evidently, single-cores performance has stalled and new processing solutions need to be studied to create the new products and services these projects target. Furthermore, apart from what the technology and the industry will deploy, the safety requirements (e.g., space and time partitioning) will also be required to be implemented, and security requirements have to be addressed to satisfy the openness and integration needs.

3.1.2.1 Support for More Performance

In order to satisfy the previously described performance requirements, alternative approaches need to be explored. Multi-core processors are the most promising alternative to the single-core architectures used in the past. Theoretically, multi-core would satisfy the new performance requirements. In addition, they are widely applicable as the current software applications can be reused.

However, in order to use multi-core processors in safety-critical solutions, a temporal and spatial partitioning between the deployed application(s) need to be ensured. Hardware solutions, like the

MMU can still be used to ensure space partitioning, but due to the interferences time partitioning is difficult to ensure, when using shared resources on multi cores. These interferences may come from the usage of processor peripherals or accelerators (e.g., Direct Memory Access (DMA)), or from the caches, buses and memories that are shared between multiple cores. Studies as [69, 70] proved that applications suffer from a slowdown in the execution time proportional to the number of processors cores due to the interferences in the usage of the shared resources (caches, buses and memory). These studies effectively render the usage of multi-cores inefficient on safety-critical solutions, unless mitigation solutions are introduced.

Hardware [71–74] and/or software [75–82] approaches have been proposed to address time partitioning on multi-core processors. However, the following requirements of safety-critical industrial products [83] demand for additional consideration to be taken into account: support for legacy applications, efficiency of the proposed software/hardware solution, robust partitioning assessment complexity, integration into an industrial process, easiness to adapt existing applications, and complexity and certifiability of the solution.

3.1.2.2 Support for Mixed-Criticality

With the advent of multi-core processors, new challenges were introduced for the integration of applications of different criticality levels. IMA solutions have been used for the integration of applications on single-core processors. In multi-core processors, new approaches need to be yet defined to ensure that for safety-critical applications time and space partitioning are respected and at the same time, low-critical applications employ the highest possible performance.

Fisher [84] introduced a commercial solution that enabled the combination of critical and non-critical applications by disabling the execution of non-critical applications when a critical application was executed in the system. In this solution the execution of non-critical applications on all available cores is disabled, only when the critical applications (executing in a single core of the processor) are not scheduled. This effectively reduces the utilization of the system, as the cores remain unused when the critical application executes.

Multiple projects (including DREAMS) and studies [85–88] have addressed (or are addressing) this topic, but few have addressed this topic while considering solutions for integrating multiple critical applications in a multi-core processor (see previous subsection in Section 3.1.2.1). Moreover, performance is not the only aspect that multi-core processors introduce, when combining critical and non-critical applications in a system. This integration brings new challenges and solutions, when considering hardware faults. For example, when a core fault occurs, it does not mean that the whole processor needs to be disregarded. However, fault-tolerance mechanisms need to be introduced in the development and execution of such systems to ensure that the safety-critical functionalities of the system are maintained in such events.

3.1.2.3 Cyber-Security

Safety-critical systems have been mostly autonomous and disconnected from non-critical systems. However with the introduction of electronics on those systems, there is now a trend to connect and furthermore integrate them, e.g., federated systems, IMA. With the new requirements added for solutions to further enhance the aircraft safety and provide new exploitation services, these previously completely autonomous systems require now to communicate to the external world. As a result, to ensure the safety of those systems, security becomes essential, for instance, the security support virtualization for safety-critical systems [89–92].

3.2 Wind-Power Domain

In the wind-power domain, there is a tremendous market push towards off-shore operation. The road to off-shore introduces new technological challenges, stringent safety requirements and new standards to comply with. In this section, the state-of-the-art solutions in the wind-power domain are discussed.

3.2.1 State-of-the-Art: Wind-Turbine Control and Supervision System

Commercial wind-turbines are governed by a control and supervision platform, which implements the following two groups of functionalities:

- Control and supervision
- Human-machine interface and communications with the Supervisory Control and Data Acquisition (SCADA)

Figure 3.1: Galileo Platform Version 5.0

Figure 3.1 represents the latest version of the Galileo platform that is currently used in the state-of-the-art wind-turbines. Galileo is a commercial hardware (industrial PC APC 910) based on an x86 dual core processor that customized at operating system and software levels. Though the Galileo platform is used for the control and supervision of the wind-turbines, it may support other real-time applications such as wind farm control. The Galileo platform requires several inputs and outputs that are connected through an EtherCAT field bus.

3.2.2 State-of-the-Art: Safety Protection System

The protection system is in charge of maintaining the wind-turbine in a safe state, by assuring that the design limits of the wind-turbine are not exceeded. The protection functions are activated as a result of a failure of the control function (running in the supervisory system) or of the effects of an internal or external failure or dangerous event. It should be activated in cases such as:

- Over-speed

- Generator overload or fault
- Excessive vibration
- Abnormal cable twist (due to nacelle rotation by yawing)

The state-of-the-art protection system is an external module integrated in the EtherCAT ring. This is the only module that checks safety data in order to decide over the safety chain. The protection system works independently from the Galileo platform, just sharing the EtherCAT bus, as shown in Figure 3.2.

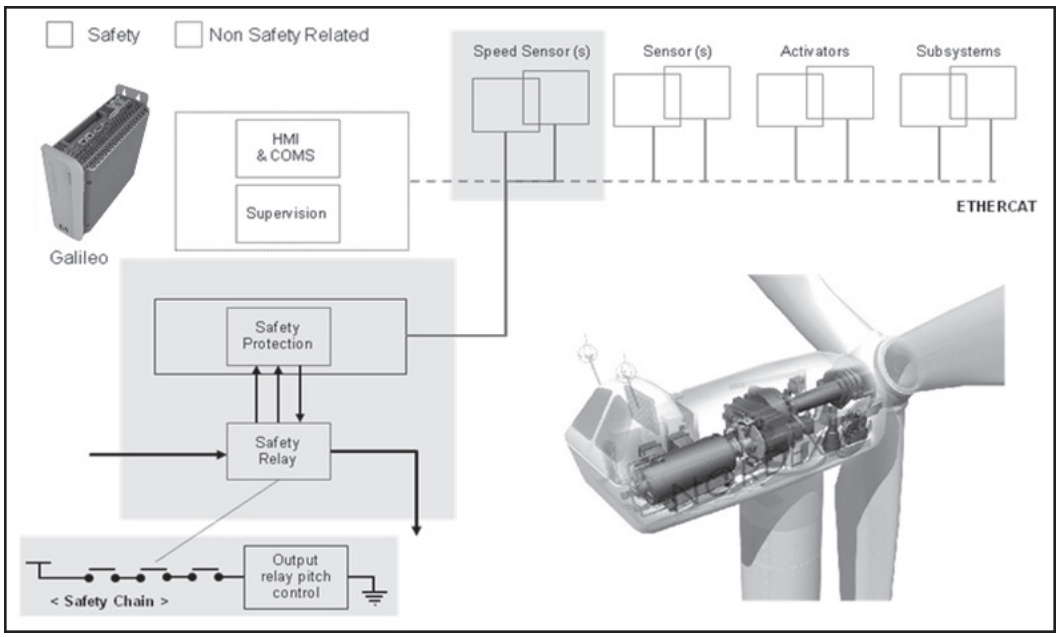

Figure 3.2: Overview of the Safety Protection System [93]

The act of checking safety data and deciding over the safety line is done in a non-redundant way. Thus, the current solution is a zero Hardware Fault Tolerance (HFT) system, which would mean that a failure in the safety protection system causes the loss of the safety function.

3.2.3 Challenges

The proposed state-of-the-art solutions for the control and supervision system and the protection system face several challenges:

3.2.3.1 Integration and Flexibility (Cost Reduction)

The protection system is implemented in an external module that is integrated in the EtherCAT ring and it lacks the flexibility to implement complex logics, as it is able to only handle digital inputs and outputs and it is mainly a commercial hardware based system.

Furthermore, there are other limitations, such as a reduced reliability and availability due to interconnection of subsystems (low integration), improvable maintainability due to the high number of functionalities being executed in different devices, operating systems not suitable for the purpose of all hosted applications, limited scalability and lack of composability (introducing more functionalities in such a complex and static architecture is not an easy task).

3.2.3.2 Mixed-Criticality

Galileo control and supervision platform is home to different kinds of processes and tasks. High criticality processes such as wind power control algorithms are run together with lower criticality processes such as external communication management. An overload in low criticality tasks could affect the performance of other processes such as the control algorithms.

3.2.3.3 Safe Communication

Galileo platform uses redundant EtherCAT communication in a ring topology to communicate with all subsystems in the wind turbine. Most of these subsystems are not safety-related and therefore the EtherCAT solution has been configured to maximize availability, but without safety concerns. Normally, Fail Safe over EtherCAT (FSoE) is used to add safety traffic over this EtherCAT network, but this protocol is not open and has some limitations imposed by the manufacturer.

Safety data is sent by an EtherCAT node to the safety protection system which is connected to the EtherCAT ring. The safety protection unit then evaluates this data and decides whether the safety line should be opened or not. Data integrity is not guaranteed in this communication channel. Therefore, decisions over the safety line can be taken based on corrupted data.

3.2.3.4 Safety Certification

The current solution is a zero HFT system and achieving SIL 3 integrity level is difficult.

3.3 Health-Care Domain

In the health-care domain there is an increasing trend towards mobile and smart solutions, to achieve potential improved care, cost savings and reduction of errors. Such solutions allow physicians to monitor patients remotely in real-time, ultimately avoiding medical errors, providing patient comfort, while also reducing treatment costs. This section offers an overview of the challenges and state-of-the-art solutions in the health-care domain.

3.3.1 State-of-the-Art Solutions

Miniaturized medical sensors detect biological signals and extract health-care data, assisting to a proliferation of services based on wearable intelligent devices [94–96]. In particular, heart-related disorders account for one of three deaths in the US in 2017 and include *atrial fibrillation* related to stroke risk and *ventricular arrhythmia* associated to cardiac arrest [97]. These disorders occur sporadically and are now treated with off-line ECG analysis based on 72-hour ambulatory Holter that has a low diagnosis rate.

Prolonged real-time ECG monitoring is needed for improved detection. However, most industrial products, e.g., AliveCor [98, 99], BodyGuardian [100], LifeMonitor [101], NowCardio [102], and PhysioMem [103], capture, process and transmit ECG data to a server for off-line analysis by specialists. In addition, research on wearable monitoring and arrhythmia diagnosis concentrates on detection capability than real-time, cf. Android [104] or iOS [99]. Android smart-phone is often used to capture, analyze and visualize ECG for alerting the patient in real time [105, 106]. Transmission of annotated ECG signals to a remote monitoring center for *off-line processing* by medical personnel is considered in [107–109], whereas ECG data packetization and TCP parameters can be considered prior to network transmission [110].

Two years ago, a lightweight mobile wearable cardiac pulse sensor (called BodyGateway, or

BGW) was developed with several micro-electromechanical sensors and an ARM Cortex M3 microcontroller supporting real-time OS. The BGW is an open version of BodyGuardian, which is attached to one of three places on the patient's chest (similar to a bandage) and allows physicians and care providers to monitor important biometric data (ECG signal, heart rate, respiration rate and physical activity level and body position) in real-time. The patch can be programmed to either stream or store and periodically transmit vital physiological data via Bluetooth over a continuous 30-day period to a host device.

The BGW can be used for remote monitoring for in- and out-hospital use cases, such as *rhythm monitoring* to understand the cardiac role of rare unexplained symptoms, *vitals monitoring* to study cardiac rhythm respiration and activity, or *long-term treatment effectiveness monitoring* to evaluate arrhythmia medication therapy. In all cases, we avoid expensive clinical trials, whether it is for daily checkups, or after heart attacks, surgery, or implants, while reducing patient concerns and improving engagement with care plans.

Recently the BGW functionality was integrated in a BodyGuardian Heart product to provide a discreet, pocket-sized wearable monitor. This solution accommodates patient mobility, enhances compliance, and streamlines data collection via wireless. In addition, a *Remote Monitoring Center* allows physicians to monitor patients' experience for individualized monitoring and care plan support. In this context, BodyGuardian Heart supports the following functionalities:

- Remote monitoring of ambulatory ECG and average heart rate for arrhythmias, including Atrial Fibrillation, Tachycardia, Bradycardia, Pause and others
- Recording and wireless transmission of periodic ECG at specified intervals for static analysis, so that physicians can access their patients' data set and review anytime, anywhere periodic cardiac event notifications (e.g., maximum and minimum heart-rate) via the web in a secure way through the PatientView and PatientFlow portals or a connected electronic medical record system

3.3.2 Challenges: Platform Security and Functionality

Digital medical information, whether stored on a computer or transmitted via wireless is vulnerable to hackers and fraudsters. In fact, there is a much higher rate of cyber attacks, e.g., as identity theft, on health-care providers and insurers. Despite rigorous security precautions and governmental rules and civil penalties, medical data is not secure, as demonstrated by recent hacks, e.g., Verizon's data breach affected 14 million patients.

The Health-care Insurance Portability and Accountability Act (HIPAA) establishes US standards for the protection of health information data, including technological safeguards for enforcing compliance. The BodyGuardian heart monitoring system complies with HIPAA Privacy and Security Rules, by supporting device-, network- and application-level security services.

A major challenge in DREAMS focuses on extending the state-of-the-art in biometric signal processing by realizing a health-care demonstrator, which exploits the advent of real-time and time-triggered technologies. More specifically, required support for a real-time diagnostic ECG arrhythmia detection application was considered, which can periodically capture and communicate in real-time asymptomatic events based on patient's vital signs to physicians, even in the presence of mixed-criticality traffic, such as infotainment. Settings of the ECG analysis algorithms can be customized to other non-fatal arrhythmias with high predictability.

In Chapter 7, key hardware/software architectural components are addressed. Furthermore, Chapter 11 focuses on the preliminary evaluation of the DREAMS health-care demonstrator solution in both in- and out-of-hospital scenarios. In addition, the health-care demonstrator supports hard real-time communication components (XtratuM hypervisor, TTEthernet router/driver).