


TOPICAL WORKSHOP ON ELECTRONICS FOR PARTICLE PHYSICS  
GEREMEAS, SARDINIA, ITALY  
1–6 OCTOBER 2023

## Using Software Mitigation Schemes to improve the availability of IoT applications in harsh radiation environment

A. Zimmaro <sup>a,b,\*</sup>, R. Ferraro,<sup>a</sup> J. Boch,<sup>b</sup> F. Saigne,<sup>b</sup> A. Masi<sup>a</sup> and S. Danzeca<sup>a</sup>

<sup>a</sup>CERN, Organisation Européenne Pour la Recherche Nucléaire,  
CH-1211 Genève, Switzerland

<sup>b</sup>IES, Université de Montpellier, CNRS,  
Montpellier, France

E-mail: [alessandro.zimmaro@cern.ch](mailto:alessandro.zimmaro@cern.ch)

**ABSTRACT:** The integration of IoT infrastructure in the context of particle accelerators promises numerous benefits (reduced costs and maintenance time, increased deployment). However, the use of microcontroller units (MCUs), typical of IoT systems, can potentially compromise future accelerator availability performances. This paper presents Software Mitigation Schemes (SMS) designed to improve the availability performance of MCU-based systems under radiation. Their effectiveness is demonstrated through a radiation test on a CERN Wireless IoT Radiation Monitoring system, also called BatMon. The results underline the IoT devices' feasibility as a viable solution for high-distribution systems in the High-Luminosity Large Hadron Collider (HL-LHC) or Future Circular Collider (FCC).

**KEYWORDS:** Accelerator Applications; Radiation damage to electronic components; Radiation damage monitoring systems; Digital electronic circuits

\*Corresponding author.

---

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Wireless IoT Monitoring system: hardware and firmware choices</b>	<b>2</b>
<b>3</b>	<b>Software Mitigation Schemes description</b>	<b>3</b>
<b>4</b>	<b>Availability performance enhancement through SMSs</b>	<b>4</b>
<b>5</b>	<b>Conclusion</b>	<b>5</b>

---

## 1 Introduction

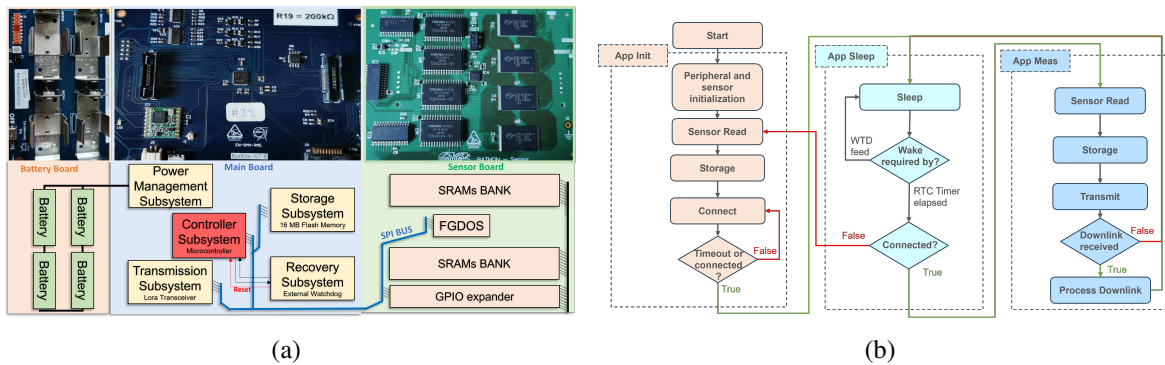
In the context of particle accelerators, the concept of the Internet of Things (IoT) is gaining increasing interest because of the benefits it could bring to this sector: a) reduce the time required to deploy and replace systems, b) higher distribution of system and sensors thanks to the no need of cables and c) cost. The different layers of the IoT, find a similarity in particle accelerators, where all the equipment, sensors, and devices are intrinsically interconnected into what is commonly known as a ‘control system’ [1]. Nevertheless, this concept also poses new challenges when applied in a radiation environment. In [2], the challenges of qualifying low-power components were shown, while in [3], the protocol that best suits the LHC environment was discussed. Another challenge that has not been yet discussed is the performance in terms of availability, defined as the operating time of the machine compared to the overall operating time [4], that an IoT device can guarantee in the accelerator environment. In this sector, the concept of machine availability is normally considered more critical than that of reliability. Unlike space, it is always possible to access the equipment and perform preventive maintenance on it, extending its lifetime indefinitely. On the other hand, when a component stops working due to SEE, it is necessary to stop the machine, identify the cause of the problem, and eventually replace or restore the faulty equipment. Increasing the availability of the machine has been the main mandate of the Radiation to Electronic (R2E) project [5], whose effort allows to reduce the number of radiation to electronic failure dump per year from 12 dump/fb<sup>-1</sup> in 2011 to <0.5 dump/fb<sup>-1</sup> in the latest years [6, 7], through introduction of protection shieldings, redesigning of equipment and introduction of mitigation scheme. For HL-LHC, the scope is to improve the machine availability to 0.1 dump/fb<sup>-1</sup>, and as a consequence, the introduction of IoT devices has to not impact this parameter. These devices, to meet the requirements of low cost, high distribution, and battery power, normally use MCU instead of FPGA, which are known to be more reliable under radiation. Existing mitigation schemes available in the literature, are not suitable for IoT devices since impact their performances. For instance, memory scrubbling algorithms to correct SEEs that occur in RAM [8, 9], such as the Single-Error Correct Double Error Detect scheme, and affect the execution of the algorithm are not an option, either because they are not available on commercial MCUs or because they can impact performance in terms of power consumption. At the firmware level, it is still possible to act by tripling the computational algorithms and correcting SEUs occurring. In [10], the authors show how the Trikaya algorithm, which consists of triplicating the subroutine constituting the algorithm with a majority voting, can reduce corrupted

calculations by five to ten times. However, as has been reported, its implementation causes an increase in execution and energy consumption, and multiple SEUs can still cause the algorithm to fail. An improved mitigation scheme is proposed in [11]. In this case, the algorithm is always triplicated, but the three subroutines are not always executed. If the output of 2 first-executed subroutines is equal, the third one is skipped. Conversely, if the two outputs do not match, it is executed, and the output is corrected by a majority voting. This approach, compared to [10], reduces execution time and power consumption, maintaining the performance of the previous work in terms of reducing damaged calculations. Automatic tools that allow the entire algorithm to be triplicated exist such as [12]. Nevertheless, those approaches are not compliant with IoT design, since they impact the performances of the devices increasing the execution algorithm time. In addition, IoT design exploits several states during operations using different peripherals: those mitigation schemes are not capable of detecting if external or internal hardware is not working as it should and recovering the system functionalities.

This paper presents new software mitigation schemes designed for IoT MCU-based design capable of improving the availability performance of the design under radiation and not impacting the system performances. Their effectiveness is demonstrated by their direct application to a Wireless IoT Radiation Monitoring system for the LHC, the BatMon [2, 3], described in section 2, and through a comparison in terms of system performance in terms of availability under radiation. This paper demonstrates that the usage of IoT in an accelerator environment will not impact the performance of availability in the accelerator, bringing all the advantages mentioned above.

## 2 Wireless IoT Monitoring system: hardware and firmware choices

Needs of high mobility and versatility, and low cost drove CERN to develop the first Wireless IoT Radiation Monitoring system at CERN. Comprising three swappable boards — Power, Mother, and Sensor — (figure 1(a)), the Mother Board, is the main layer of the system since its purpose is to manage the system operation. Within it is possible to distinguish different subsystems. The controller subsystem is the heart of the entire design and is based on a microcontroller (MCU), choice based to be compliant with the previously mentioned requirements. The Transmission subsystem is used for wireless communication and is a LoRa transceiver. A recovery subsystem, made of an External Watchdog (Ext WTD) is used to recover systems from Single Event Functional Interruptions (SEFIs), as in [13]. The radiation monitoring sensor-board, includes a Floating Gate Dosimeter (TID) and




**Figure 1.** Wireless IoT board and general hardware architecture (figure 1(a)) and firmware flowchart (figure 1(b)).

calibrated SRAMs sensitive for Thermal Neutron (ThN) and High Energy Hadrons (HEH) fluences [2]. Due to the high number of required GPIOs, SPI-controllable GPIO expanders are used by the MCU. The system uses all Commercial of the Shelf (COTS) low-power and qualified components [2]. A Finite State Machine of three states is used to control the system (figure 1(b)). After start-up (setting up of the different clocks and GPIOs), the system initializes the sensors and peripherals required for operation. Once completed, the sensors are read, the measurements stored in the storage subsystem, and try to connect. For the system to take account of any network unavailability, a timeout is defined after which the system stops trying to connect and enters sleep mode (App sleep). In case it connects, the previous measurements are transmitted and App sleep is reached. In that mode, the system reduces energy consumption by disabling peripheral devices not required during sleep operation or setting them to low-power mode. A Real Time Clock (RTC) configured with a certain timeout (Duty Cycle Time), is used to wake the system at the end of this period. If it is not connected, the system measures the sensor again and attempts to reconnect (App Init). If it is connected, the state App Meas is reached, the sensor read, the measurement transmitted and it goes back to App Sleep. Once connected, the system no longer enters the App Init phase. During all phases, the MCU feeds the Ext WTD via an external interrupt. By changing the App sleep duration, it is possible to increase the system lifetime [2].

As shown, normally IoT system performs several operations by exploiting different peripherals during its functioning. The triplication of such operations would increase the computing time of the device and may not mitigate any SEE. Furthermore, their effect on system functionality, unless caused by a CPU Crash, may be invisible to the Ext WTD. Software Mitigation Schemes (SMSs), able to detect system malfunctions and restore functionality are required to have reliable systems.

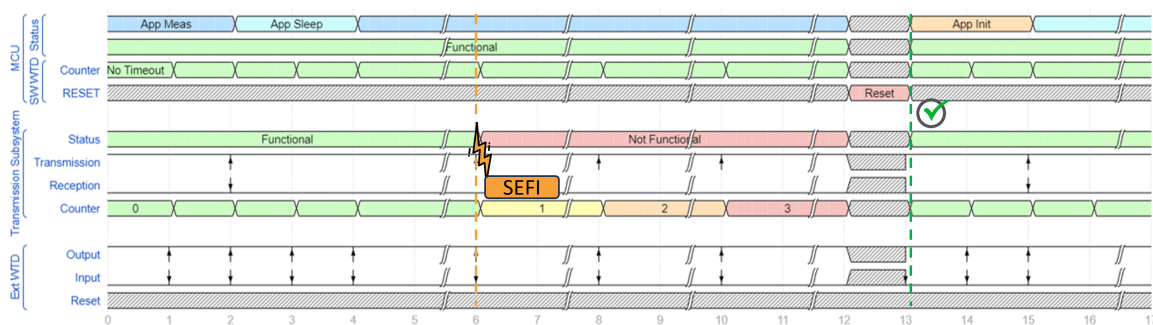
### 3 Software Mitigation Schemes description

As described in section 2, the Wireless IoT design embeds a recovery subsystem made of an Ext WTD. The internal watchdog peripheral, when available, can be used to detect failures invisible to the Ext WTD. In the case of the Wireless IoT this peripheral was not usable since its lifetime was limited at a TID of around 100 Gy. As an alternative, an Internal Software Watchdog (SW WTD) was designed. It uses the RTC, checking the time elapsed at each wake triggered by the Ext WTD. If the Software timeout defined is reached, the MCU is stacked in one of the three operational states (i.e. infinite while loop, no wake from sleep). This type of error can be invisible to the Ext WTD since the MCU may be always able to feed it. The working principle of the SW WTD is depicted in figure 2.



**Figure 2.** The SW WTD recovery working principle is depicted. At tick 7 a failure occurs which the Ext WTD cannot detect because the MCU continues to reply. Due to the failure, the MCU does not wake up and continue sleeping. At tick 10, when the watchdog wakes up the system since the timeout has been reached, the SW WTD recognizes that the system is not functional and self-resets.

The transmission subsystem, used to perform wireless communication, has low observability from a failure point of view. The MCU can only control and configure it via SPI but it cannot monitor its status. It was observed during system qualification a failure consisting of stop transmission capability. The chip signature and configuration were still readable when the failure occurred. To detect this failure, a specific software mitigation scheme was designed. The LoRa Confirmed Uplink (CU) feature is used to check if the transceiver is really transmitting. A CU from the Network Server (NS) is requested for every configurable number of transmissions. When not received for three consecutive times, it means that the transceiver is not working as it should and is reset. This number of consecutive failed receptions has been set to 3. The choice of 2 would have saved time (one Duty Cycle Time that corresponds to 5 minutes in the example described), but it was observed that the possibility of having 2 consecutive packets not received by the NS was not negligible. A flowchart describing the mitigation scheme is depicted in figure 3.



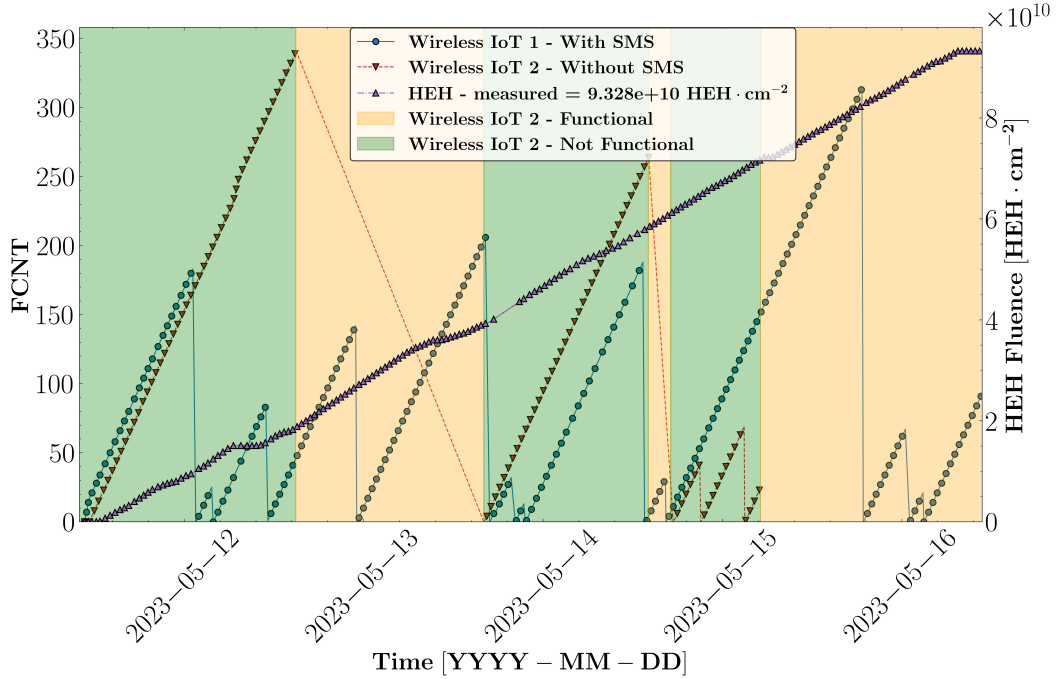
**Figure 3.** The mitigation scheme used to detect a failure of the LoRa transceiver is depicted. As visible, the failure of the transceiver is invisible both to the Ext and SW WTDs. For every configurable number of transmissions, the MCU requests a CU and updates a counter. If the confirmation is received (Tick 2), the transceiver is considered functional. If the confirmation is not received as in the occurrence of Tick 6, a counter is updated and a confirmation is requested at the next uplink. In case it is not received 3 consecutive times, the MCU resets itself and the Transceiver.

In addition to these mitigation schemes, a further protection mechanism against MCU misconfigurations was implemented. Management function via attribute instruction is assigned to all handler function definitions in the code and provides a safeguard against wrong configurations that could result from SEEs. When the latter produces a setting change of a handler in the MCU, the above mentioned management function is executed which reports the error and resets the device. It was observed that such type of error was also invisible to the Ext WTD.

#### 4 Availability performance enhancement through SMSs

To verify the impact on system operation in terms of availability, it was decided to test two Wireless IoTs at CHARM: one using SMS and one using only Ext WTD as mitigation. CHARM is a mixed-field irradiation facility designed to qualify entire systems within a realistic field that is fully representative of the mixed-field environment of the high-energy accelerator [14]. In figure 4, the LoRa Frame Counters (FCNT) (a counter that is incremented by one with each transmission) of the two devices, recorded during the tests in CHARM, are shown. As it is visible from figure 4, the Wireless IoT without the SMS (Wireless IoT 2) has a longer period of downtime respect the Wireless IoT 1 (3517 vs 187 minutes).

The downtime/uptime is reduced from 48.82 to 2.58%. The downtime of 187 minutes cumulated during the run does not represent a limitation for the application in terms of availability requirements. The fluence cumulated in CHARM corresponds to 1.87 years of operation in the Dispersion Suppressor (DS) of HL-LHC ( $5 \cdot 10^{10} \text{ cm}^{-2}$ ) [15]. Considering the downtime observed in CHARM and the time to cumulate such radiation level in DS, the expected unavailability per year will be 0.01%.



**Figure 4.** The LoRa FCNT is depicted for the two Wireless IoT tested in CHARM. The Wireless IoT without the SMS (Wireless IoT 2 in red), was affected by SEFIs not detected by the only Ext WTD. During these malfunctioning periods (in Orange), the device was not able to transmit or was stacked in a while loop or sleep mode. On the other hand, when the SMS was implemented (Wireless IoT 1 in blue), the device run without interruption throughout the test, with only a few intervals without any data due to error detection and recovery times.

## 5 Conclusion

In this paper, Software Mitigation Schemes for MCU-based designs going through different operational states, typical of IoT devices, have been presented. Their impact on system availability under radiation was validated in CHARM. Considering CERN’s availability requirements for critical systems at the LHC (99.54%) [16], an IoT application with Wireless IoT-like performance, i.e. an expected availability in the DS area of 99.99%, can foresee 45 systems in this area while respecting the LHC constraints. Since application availability depends on the annual fluence, it will be higher in low radiation areas (i.e. shielded areas), allowing more devices to be used. These results demonstrate that IoT devices can be a viable solution for critical high-distribution systems in the future HL-LHC or FCC. Furthermore, due to its modularity and versatility, the Wireless IoT can also be used for more critical applications, such as wireless remote control of equipment.

## References

- [1] S. Deghaye and E. Fortescue-Beck, *Introduction to the BE-CO Control System*, [CERN-ACC-NOTE-2020-0069](#) (2020).
- [2] A. Zimmaro et al., *Testing and Validation Methodology for a Radiation Monitoring System for Electronics in Particle Accelerators*, *IEEE Trans. Nucl. Sci.* **69** (2022) 1642.
- [3] S. Danzeca et al., *Wireless IoT in Particle Accelerators: A Proof of Concept with the IoT Radiation Monitor at CERN*, *JACoW IPAC2022* (2022) 772.
- [4] A. Apollonio et al., *Lessons Learnt from the 2016 LHC Run and Prospects for HL-LHC Availability*, *JACoW IPAC2017* (2017) TUPVA006 [CERN-ACC-2017-145].
- [5] *R2e mitigation project*, <https://www.cern.ch/r2e>.
- [6] M. Brugger, *R2E and availability*, *CERN Yellow Rep.* **2** (2015) 149.
- [7] Y. Aguiar et al., *Radiation to Electronics Impact on CERN LHC Operation: Run 2 Overview and HL-LHC Outlook*, *JACoW IPAC2021* (2021) MOPAB013.
- [8] H. Quinn, *Challenges in testing complex systems*, *IEEE Trans. Nucl. Sci.* **61** (2014) 766.
- [9] G. Tsiliogiannis et al., *Radiation Effects on Deep Submicrometer SRAM-based FPGAs under the CERN Mixed-Field Radiation Environment*, *IEEE Trans. Nucl. Sci.* **65** (2018) 1511.
- [10] H. Quinn, Z. Baker, T. Fairbanks, J.L. Tripp and G. Duran, *Software resilience and the effectiveness of software mitigation in microcontrollers*, *IEEE Trans. Nucl. Sci.* **62** (2015) 2532.
- [11] H. Quinn, Z. Baker, T. Fairbanks, J.L. Tripp and G. Duran, *Robust duplication with comparison methods in microcontrollers*, *IEEE Trans. Nucl. Sci.* **64** (2017) 338.
- [12] M. Bohman et al., *Microcontroller compiler-assisted software fault tolerance*, *IEEE Trans. Nucl. Sci.* **66** (2019) 223.
- [13] D.A. Santos et al., *Characterization of a risc-v system-on-chip under neutron radiation*, in the proceedings of the *2021 16th International Conference on Design and Technology of Integrated Systems in Nanoscale Era (DTIS)*, Montpellier, France (2021), p. 1–6, [DOI:10.1109/DTIS53253.2021.9505054].
- [14] J. Mekki et al., *CHARM: A Mixed Field Facility at CERN for Radiation Tests in Ground, Atmospheric, Space and Accelerator Representative Environments*, *IEEE Trans. Nucl. Sci.* **63** (2016) 2106.
- [15] R. García Alía et al., *LHC and HL-LHC: Present and future radiation environment in the high-luminosity collision points and RHA implications*, *IEEE Trans. Nucl. Sci.* **65** (2018) 448.
- [16] T. Cartier-Michaud et al., *Data-Driven Risk Matrices for CERN's Accelerators*, *JACoW IPAC2021* (2021) 2260.