POLITECNICO DI TORINO
Repository ISTITUZIONALE

Evaluating Reliability against SEE of Embedded Systems: A Comparison of RTOS and Bare-metal Approaches

Original
Evaluating Reliability against SEE of Embedded Systems: A Comparison of RTOS and Bare-metal Approaches / DE SIO, C., Azimi, S., Sterpone, L.. - In: MICROELECTRONICS RELIABILITY. - ISSN 0026-2714. - 150:(2023). [10.1016/j.microrel.2023.115124]

Availability:
This version is available at: 11583/2981728 since: 2023-09-06T12:28:58Z

Publisher:
Elsevier

Published
DOI:10.1016/j.microrel.2023.115124

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

16 October 2026

Evaluating Reliability against SEE of Embedded Systems: A Comparison of RTOS and Bare-metal Approaches

C. De Sio[a*], S. Azimi[a], L. Sterpone[a]

[a] *Department of Computer and Control Engineering, Politecnico di Torino, Turin, Italy*

Abstract

Embedded processors are widely used in critical applications such as space missions, where reliability is mandatory for the success of missions. Due to the increasing application complexity, the number of systems using Real-Time Operating Systems (RTOSs) is quickly growing to manage the execution of multiple applications and meet timing constraints. However, whether operating systems or bare-metal applications provide higher reliability is still being determined. We present a comprehensive reliability analysis of software applications running on a device with bare-metal and FreeRTOS against the same faults based on fault models derived from a proton test. Additionally, the FreeRTOS system has been evaluated with a set of software applications dedicated to evaluating specific RTOS functions, providing an additional evaluation for operations crucial for a real-time operating system.

1. Introduction

Embedded processors have become increasingly popular in mission-critical applications such as space missions, where reliability is paramount.

The impact of soft errors on reliability can be significant in these systems, making it essential to ensure that the applications used are reliable and robust. Moreover, the continuous downscaling of transistors and operating voltages has led to more performant devices. Such devices are appealing for high-performance mission-critical applications, such as space missions. However, smaller transistors dimensions, operating voltages, and higher frequency made them more vulnerable to soft errors, which is a primary concern for systems deployed in harsh environments, such as space, where the exposure to ionizing radiation is a source of malfunctions in the device [1]. As the complexity of tasks that embedded systems must perform continues to increase, the bare-metal approach has decreased, leading to migration towards adopting Real-Time Operating Systems (RTOSs). These systems provide an efficient solution for meeting stringent real-time requirements, particularly in safety-critical applications that require the management of the execution of multiple critical applications on the same platform [2]. Despite the widespread use of embedded processors, the robustness of software applications running on systems with RTOS compared to bare-metal has yet to be thoroughly investigated. This leads to a broad question of whether these two platforms' reliability differences exist when they run applications in safety-critical missions.

1.1 Main Contributions

This paper is dedicated to performing an accurate and comprehensive reliability analysis of two developed platforms running the same applications in the presence of the same fault models, relying on RTOS and bare-metal, respectively. The paper presents two main contributions. Firstly, a detailed analysis of the fault model occurring in the on-chip SRAM memory of an ARM Cortex-A9 embedded processor during a proton test is presented. The observed events are used to propose a set of fault models for realistically emulating radiation-induced soft errors. These fault models provide a model for radiation-induced errors observed in the on-chip memory from the processor side. Secondly, we proposed two reliability analyses for two platforms running on ARM Cortex-A9 embedded processor of a Zynq-7020 system-on-chip. FreeRTOS and bare-metal platforms are evaluated using the same software applications suite. Finally, since RTOS systems are characterized by additional features that are missing in bare-metal, an additional analysis based on a suite of benchmark applications is provided for exploring the robustness of specific RTOS functions.

* Corresponding author. corrado.desio@polito.it
Tel: +39 (011) 0907012

2. Related Works

Several works have investigated the impact of soft errors on processor systems using different techniques [3]. During accelerated radiation testing, an embedded device is exposed to a high flux of radiation while the application is running, simulating years of functioning in a radiation environment such as space [4].
Even if accelerated radiation testing provides the closest results to the actual case scenario, it has a high cost of money, time, and expertise, making it unsuitable for the early stages of the design development flow. Therefore, other works are dedicated to developing alternative techniques, such as simulation and emulation environments, for assessing the reliability of bare-metal applications running on microprocessors. Fault Injection techniques, as one of the most common emulation approach, is widely exploited to evaluate the impact of SEEs on embedded processors [5][6]. Simulation-based fault injection methodologies are developed for emulating fault by injecting faults in memory resources, CPU registers, and communication infrastructure [7], while the impact of soft errors on the operations of a microprocessor-based architecture by injecting random SEU at a random time is investigated in [8].

However, on the other hand, the increasing complexity of tasks required for embedded systems has led to the rise in the adoption of Real-Time Operating Systems (RTOSs), which provide an efficient solution for meeting stringent real-time requirements. Therefore, with the emergence of RTOS, some works have also investigated software-level techniques for evaluating the sensitivity of software executing on embedded processors with an operating system to soft errors [9]. The vulnerability of FreeRTOS has been evaluated through a software-based fault injection method that targets the most relevant variables and data structure [10], while the authors in [11] developed a workflow for automatic fault injection into program and data memory. Common approaches are based on modifying the operating system kernel or altering the memory content [12]. However, software application-level methods abstract from underlying hardware architecture when considering the impact of faults on the operating system's functionality. Therefore, other approaches are based on the simulation of the hardware description of the embedded processing system, which allows the injection of upsets into registers and hidden elements at any time [13][14].
Although numerous research studies have focused on assessing the dependability of embedded microprocessors using both bare-metal and FreeRTOS approaches, none have specifically compared these approaches regarding their susceptibility to single-event effects (SEE) caused by architectural faults. Designers, particularly those working on space applications, often find themselves debating whether to use an operating system or a bare-metal application to achieve higher reliability. This paper aims to fill this gap by conducting a first direct comparison of identical applications running on embedded systems using bare-metal, and RTOS approaches, taking into consideration the fault models derived from a conducted proton radiation test experiment.

3. Fault Model Resulting from Proton Testing

We performed a proton radiation test at Switzerland's Paul Scherrer Institute (PSI) proton facility. A Zynq-7020 device has been irradiated with proton beams with energies between 29 and 200 MeV. The content of the on-chip 256 Kb SRAM memory has been continuously monitored through a software routine running on an ARM Cortex A9 during the experiment. The software test routine is executed on the processor system, reading and writing the memory content. The testing routine writes new values in the memory and checks if the value written during the previous test loop has been corrupted. It also verifies that the current value has been written correctly and can be read correctly. When an erroneous value is detected, it is notified to a host computer connected through a serial connection. The software routine can identify both Single Event Functional Interruption (SEFI) errors (e.g., a memory cell cannot be written or read correctly anymore) and soft errors, such as Single Event Upset (SEU) and Multiple Cell Upset (MCU). We identified the following fault models and

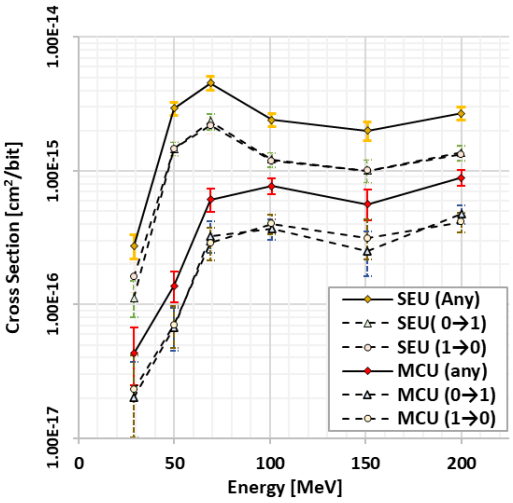

Fig. 1 SEUs and MCUs Cross-Sections

their cross-sections, reported in Figure 1, that have been adopted in the fault injection campaigns:

Single Event Upset (SEU) is the most commonly observed event during radiation experiments. An SEU produces a change in the value stored in a memory cell due to a bitflip, leading to data corruption in memory. Figure 1 displays the SEU cross-section for various proton energy experiments, while radiation-induced transitions from 0 to 1 and 1 to 0 have also been examined, revealing comparable ratios and distributions.

Multiple Cell Upset (MCU) is a group of SEUs that happen simultaneously, usually due to a single event. Indeed, a single particle can be the source of multiple upsets by exciting more than one logic cell while traversing or corrupting the memory control signals, resulting in the corruption of numerous cells. In particular, a Multiple Bit Upset (MBU) is an MCU affecting two bits of the same logical word. Interestingly, no MBU has been observed. However, detected MCUs presented recurrent patterns. The MCUs always affected the same significant bits of equally distanced logic memory words. This characteristic is likely related to the specific characteristics of the layout and architecture of the memory under test that differ from the logical organization. The distribution of the parameters, aggregating the characteristics of events occurring at all the energies, of MCUs events are reported in Table 1 and Table 2. Additional details on the proton test experiments are reported in [15].

Table 1. Normalized Occurrences of Address Offset in SEMUs

Offset in Memory between bitflips	Normalized Occurrence
128	0.61
4	0.12
124	0.06
132	0.03
others	Less than 0.01 each

Table 2. Normalized Occurrences of Affected Bits in SEMUs

Number of Multiple Upset	Normalized Occurrence
2	0.65
3	0.20
4	0.08
5	0.03
Others [6;15]	Less than 0.01 each

3. The Reliability Analysis Environment

An analysis environment has been developed to assess and compare the reliability of the two platforms against the fault models observed during the proton test experiment. The developed framework allows emulating the faults affecting the on-chip memory of the system-on-chip. Since the difference between the two platforms is only in the software stack, emulating the fault at a lower level provides a fair comparison between the two systems. In particular, the same (i.e., same fault model and location) faults are emulated and evaluated on both systems. The experiment manager runs on a host computer connected to the embedded platform through a serial connection. Using the JTAG interface of the SoC, the experiment manager can run the application on the target device, and stop the execution, manipulating the memory of the SoC to emulate the fault model, resuming the execution, and collecting the results. The reliability analysis environment, depicted in Figure 2, includes two stages. In the first stage, the system runs a version of the application under test without faults to collect data for instrumenting the fault injection process. During this stage, the environment infers information such as the application's memory footprint, average execution time, and executed instructions. This stage is fully automated and generates the information to be used in the second stage. The fault injection process flow is performed as follows: firstly, a fault injection location and time are generated. Fault location is generated among the SRAM memory of the device. Fault injection time is then generated. In this context, fault injection time refers to the point in time during the execution of the program when a fault is intentionally injected into the system. Selecting the specific point in the program's execution when the fault will be injected is done by choosing an instruction from the list of executed instructions retrieved during the instrumentation stage. By injecting faults at different times, it is possible to simulate different scenarios and observe how the system responds. Secondly, the application under test is executed in the FreeRTOS and bare-metal versions. Once the application runs, it is stopped at the specific instruction selected as the fault injection time. This is done using the debugging mode, which allows for precise control over the execution of the program. At this point, the fault is emulated at the memory level. This means that the fault is emulated into the system by altering the contents of memory in a specific way that mimics the target fault model. During the fault injection process, monitoring the system for any issues that may arise due to the injected fault is essential. To do this, the experiment manager uses timers instrumented during the process's first stage. These timers can detect system halt or endless loops, indicating that the fault has caused the program to behave unexpectedly or crash.

4. Experimental Analysis

The illustrated framework and methodology have been used for two reliability analyses. The former is dedicated to comparing the robustness of systems based on bare-metal and FreeRTOS software stacks against fault models observed in the on-chip memory of an embedded processor during proton testing by using the same suite of software benchmarks and injected faults. The latter analysis evaluates the robustness of a different suite of software applications dedicated to specific features of the RTOSs, such as task communication and scheduling.

4.1 Software Systems

We used two software suites in this section's reliability analyses. The first suite of software benchmarks, called *general-purpose software applications*, consists of four software applications. The four software are:

- *qsort:* a quick sort algorithm used for sorting arrays of data
- *matmul:* mathematical operations on matrices,
- *basicmath*: a set of basic mathematical functions, including arithmetic, trigonometric, and logarithmic functions.
- *dhrystone:* it is a computing benchmark that performs string processing operations.

Table 3 provides information on the memory footprint and the nominal execution time of these applications.

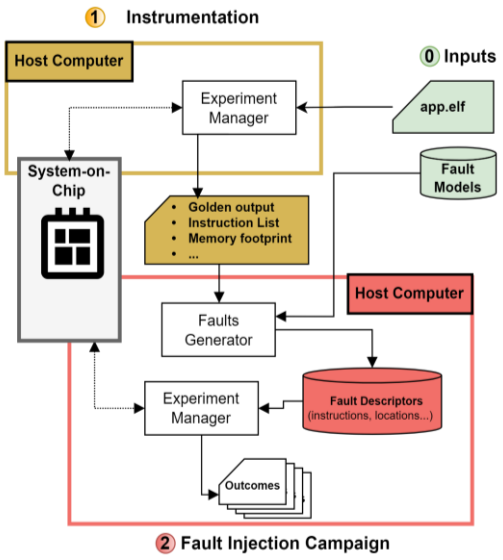

Fig. 2. Fault Injection Framework

Table 3. Characteristics of General Purpose Applications Suite

Application	Platform	Memory footprint (KByte)	Nominal execution time (ms)
qsort	bare-metal	7,932	46.45
qsort	FreeRTOS	68,796	1,049.93
matmul	bare-metal	37,580	45.67
matmul	FreeRTOS	75,348	1,047.72
basicmath	bare-metal	37,068	120.49
basicmath	FreeRTOS	74,956	1,116.15
dhrystone	bare-metal	43,636	194.69
dhrystone	FreeRTOS	78,940	1,183.18

The software applications of this suite have been implemented both for bare-metal and FreeRTOS platforms. In particular, the FreeRTOS version of each application is coded to instantiate three copies of the same task that runs concurrently on the processor. Since bare-metal systems do not support the concurrency of tasks, the bare-metal version executes the task in sequence three times through function calls to the same procedure, which consequently share the same code section, similar to what happens when multiple instances of the same task are instantiated in the FreeRTOS system.

The second suite of software benchmarks, referred to as *RTOS software applications*, consists of five software applications extracted from the Rhealstone benchmark applications suite [16]. The Rhealstone benchmark suite consists of software applications aiming to evaluate operations that are critical in a real-time operating system. In particular, the software applications are:

- *task switching*: it performs synchronous and non-preemptive task switching.
- *task preempting*: it switches tasks due to an event trigger.
- *semaphore operations*: it performs semaphore operations to support mutual exclusion between two tasks.
- *deadlock breaking*: it resolves deadlock conditions by high-priority tasks preempting a low-priority task that acquired a needed resource.
- *task communication*: it makes two tasks exchange a message.

Since these applications are dedicated to evaluating the robustness of specific features offered by RTOS systems, they have been evaluated only for the FreeRTOS platform.

Table 4 provides information on the memory footprint and the nominal execution time of these applications.

Table 4. Characteristics of ROTS Applications Suite

Application	Platform	Memory footprint (KByte)	Nominal execution time (ms)
task switching	FreeRTOS	67,348	98.68
task preempting	FreeRTOS	67,412	96.05
semaphore operations	FreeRTOS	67,668	193.06
deadlock breaking	FreeRTOS	67,859	209.21
task communication	FreeRTOS	67,548	204.67

All the reported software applications have been evaluated singularly in dedicated fault injection campaigns. All the software applications have been compiled using gcc with the -O2 optimization.
Additionally, on-board DRAM memory is not used by the hardware platform, limiting the memory space for the application (e.g., heap, stack, data, instruction, and so on) to the on-chip SRAM memory.
The memory footprint column reported in Table 3 and the same column in Table 4 does not include the heap and stack. The same stack size was used for all bare-metal applications and set to 14,336 bytes, while the heap size was set to 8,192 bytes.
The platform based on FreeRTOS was a FreeRTOS 10 version provided by ARM-Xilinx to be implemented using the Xilinx Vitis IDE v2022.1.0 in a Zynq-7020. We used the standard configuration options provided by the vendor for all the software benchmarks based on FreeRTOS. They include full support for counting semaphores and mutex, checkings for stack overflow detection enabled with methodology 2, a total heap size of 65,536 bytes, and a minimum stack size of 200 words, which is used both as minimal stack size and as stack size allocated to each FreeRTOS task [17].

4.2 Fault Injection Campaigns

We performed a dedicated fault injection campaign based on the two proposed fault models for each software application. In order to compare the bare-metal and FreeRTOS in the fairest way possible for the *general-purpose applications* suite, we evaluated them with the same fault models and locations. However, since the FreeRTOS and bare-metal versions of the application execute different instructions due to different software stacks, it is impossible to inject faults at precisely the same moment (i.e., executed instruction). The fault injection time is generated when comparing the same application on different platforms in order to emulate faults that occur at similar execution times during execution. Each fault injection campaign consisted of 10,000 experiments. During each experiment, a single fault model at a time was injected into the SRAM on-chip memory. The presented reliability analysis is based on SEU and MCU fault models emulated in the on-chip SRAM memory of the actual device, but we want to emphasize that faults in the SRAM on-chip memory are only a part of the SoC that is sensitive to radiation events. For instance, due to its high performance and minimal sizes, cache memory has the downside of being extremely sensitive to SEUs, and the choice to disable or not is still debated and based on the specific applications and their reliability and real-time constraints [18][19]. Additionally, cache memory are are also a source of unpredictability in the system that can increase the complexity of hard real-time systems. Other memories can also be used with embedded processors, such as external DDR memories. However, we chose to focus our analysis on on-chip memory since it is integrated with the SoC itself; it is the biggest on-chip memory space (e.g., compared to register files and caches) and is particularly sensitive to SEUs and MBUs.

4.3 Results Classification

The effects of the injected faults are categorized into four groups accordingly to observed impacts on the system. We identified the following categories:

- *Masked:* the fault did not visibly affect program execution. The program results are correct.
- *Silent Data Corruption (SDC):* the fault produced a corruption of the program output.
- *Crash:* the fault produced a system failure, causing the system to stop functioning. In this case, part of the output was generated before the systems stopped working.
- *Startup Failure*: the fault prevents the application from emitting any output due to an early crash or failing boot.

To clarify further the difference between a *Crash* and a *Startup Failure*, software applications have been coded to output a signature when the application under test starts to execute. A fault is classified as causing a *Startup Failure* when no output, including the starting signature of the program, is generated. Both *Crash* and *Startup Failure* cause the system to halt due to various reasons, such as endless loops or unhandled exceptions.

5. Experimental Results

The results of a first reliability analysis dedicated to evaluating the *general-purpose application suite* implemented in bare-metal and FreeRTOS platforms against the SEUs and MCUs are presented. Additionally, we present the result of a second

reliability analysis dedicated to features typical of RTOS. This dedicated reliability analysis has been carried out only for the FreeRTOS platform and evaluated against SEUs and MCUs fault models.
Reliability analyses have been conducted using statistical fault injection. We carried out fault injection campaigns of 10,000 singularly-evaluated fault injections. In accordance with [20], it allows us to reach a confidence interval of 95% with less than 1% of the margin of error of the measured error rate values. SEUs and MBUs fault models resulting from proton testing have been emulated into the on-chip SRAM memory of the embedded system.

5.1 Baremetal and FreeRTOS comparison analysis

Error rates due to SEUs and MCUs affecting the software applications of the *general-purpose application suite* running on both platforms are presented in Fig. 3 and Fig.4, respectively.
The two analyses produced similar reliability results for the evaluated applications. The resulting error rates against these fault models. Robustness comparisons among software are the same for both fault models, and the error rates vary only marginally between bare-metal or FreeRTOS based on the specific application. As a result, choosing between bare-metal or FreeRTOS can lead to slightly more robust software based on the specific application, but robustness can be considered comparable without significant variations. Since the marginal variation of the error rate when using bare metal or FreeRTOS is very small, the choice between the two mainly depends on other factors, such as more or less strong real-time requirements. However, it is interesting to notice that the distribution of the type of errors presents a pronounced difference for SEUs, which is even more marked for the MCU fault model. As shown in Fig. 5 and Fig. 6, while the error rate is similar, FreeRTOS show a significantly higher percentage of execution flow error, such as *Startup Failures and Crashes*. This is likely due to the higher complexity of the operating system layer introduced by FreeRTOS in the software stack. Differently, bare-

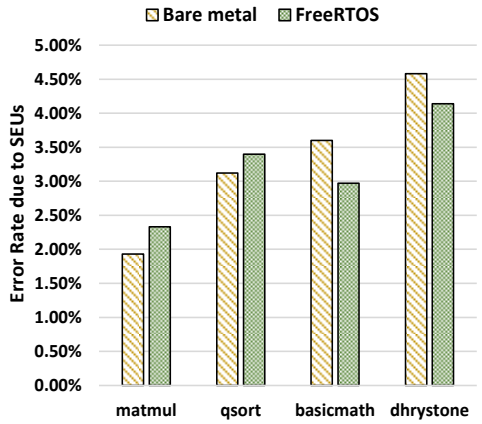

Fig. 3. Error Rate resulting from SEUs fault model for *general-purpose applications* benchmarks.

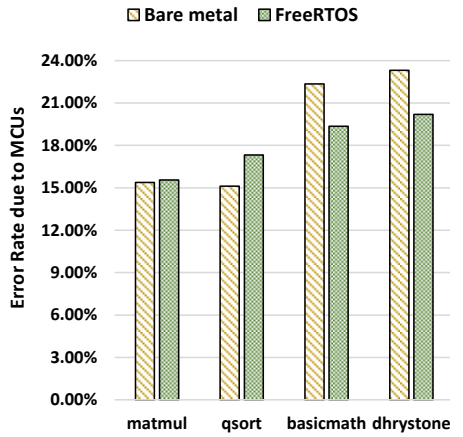

Fig. 4. Error Rate resulting from MCUs fault model

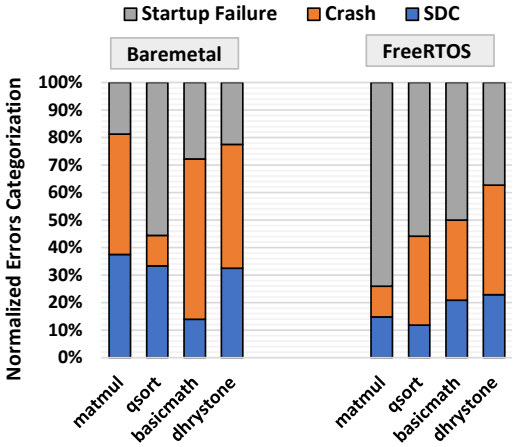

Fig. 5. Normalized Error Categorization for SEUs fault model for *general-purpose applications*

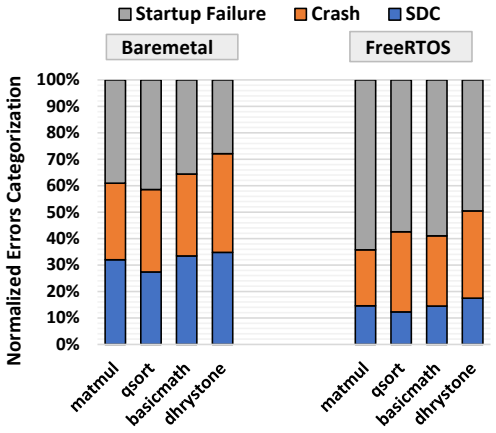

Fig. 6. Normalized Error Categorization for MCUs fault model for *general-purpose applications*.

metal is more prone to SDC errors. As a result, bare metal could be considered more suitable for systems where high availability is essential and erroneous results can be tolerated. On the other hand, SDCs are less common in FreeRTOS, which is a valuable feature since there is no advisory on the system's misbehavior in this type of error. However, since availability is an essential metric for a real-time system, this analysis raises the question of whether FreeRTOS can provide a reasonable level of availability while keeping low SDC when operating in radiation environments.

5.2 FreeRTOS Functionality Analysis

The error rates resulting from reliability analysis against SEUs and MCUs of the *RTOS benchmark suite* are reported in Fig. 7 and Fig. 8, respectively. Variations of error rates appear to be less marked among various applications, and also, the error categorization reported in Fig. 9 and Fig. 10 are very similar among the software evaluated. Due to the characteristics of the benchmark under test, SDC occurred much less compared to general-purpose applications, while errors due to control flow, such as *Crash* and *Startup Failure,* are more common. This analysis supports the idea that RTOS functionality is more prone to control flow errors, especially compared to bare-metal applications where the operating-system layer introduces much less complexity. All RTOS functionalities evaluated in this analysis seem to be characterized by a similar error rate and error categorization distribution.

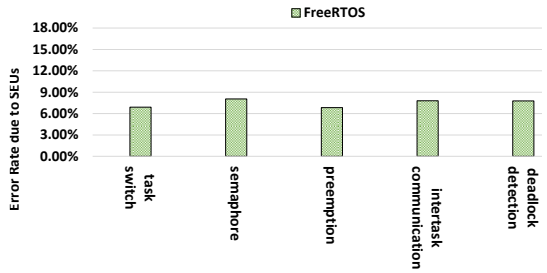

Fig. 7. Error Rate resulting from SEUs fault model for *RTOS software applications* benchmarks.

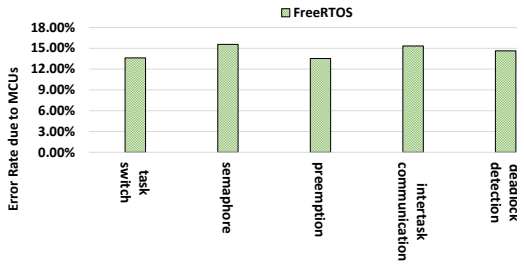

Fig. 8. Error Rate resulting from MCUs fault model for *RTOS software applications* benchmarks.

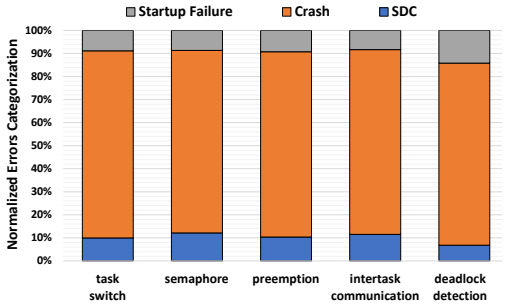

Fig. 9. Normalized Error Categorization for SEUs fault model for *RTOS software applications*.

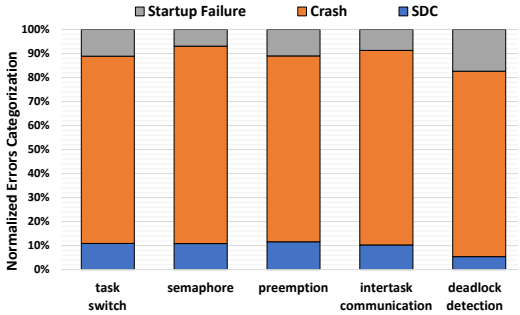

Fig. 10. Normalized Error Categorization for MCUs fault model for *RTOS software applications*.

6. Conclusions

We proposed a reliability comparison of software running in FreeRTOS and bare metal using realistic fault models of radiation-induced soft errors affecting the on-chip SRAM memory of an ARM Cortex-A9 embedded processor, such as SEUs and MCUs. Even if characterized by a similar error rate, the experimental results highlighted the different sensitivity of the two approaches to SDCs and control flow errors, which should be considered carefully when defining the software platform for real-time safety-critical applications. A second analysis dedicated to RTOS functionalities confirmed that features offered by RTOS are particularly prone to control flow errors compared to SDC.

References

[1] S. Azimi, et al., “Analysis of Single Event Effects on Embedded Processor” in MPDI Electronics, Vol 10, 2021, DOI; 10.3390/electronics10243160.

[2] A. Portaluri, et al., “Exploring the Impact of Soft Errors on the Reliability of Real-Time Embedded Operating Systems”, in MPDI Electronics, Vol 12, 2023, DOI: 10.3390/electronics12010169.

[3] C. De Sio, et al., “SEU Evaluation of Hardened-by-Replication Software in RISC-V Soft Processor”, in IEEE International Symposium on Defect and Fault

Tolerance in VLSI and Nanotechnology Systems (DFT), 2021.
[4] S. M. Guertin, et al., "Radiation Specification and Testing of Heterogenous Microprocessor SOCs," 2019 19th European Conference on Radiation and Its Effects on Components and Systems (RADECS), Montpellier, France, 2019, pp. 1-7, doi: 10.1109/RADECS47380.2019.9745708.
[5] H. Cho, S. Mirkhani, C. -Y. Cher, J. A. Abraham and S. Mitra, "Quantitative evaluation of soft error injection techniques for robust system design," 2013 50th ACM/EDAC/IEEE Design Automation Conference (DAC), Austin, TX, USA, 2013, pp. 1-10.
[6] C. De Sio, S. Azimi, A. Portaluri and L. Sterpone, "SEU Evaluation of Hardened-by-Replication Software in RISC- V Soft Processor," 2021 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT), Athens, Greece, 2021, pp. 1-6, doi: 10.1109/DFT52944.2021.9568342.
[7] F. Rosa, F. Kastensmidt, R. Reis and L. Ost, "A fast and scalable fault injection framework to evaluate multi/many-core soft error reliability," 2015 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFTS), Amherst, MA, USA, 2015, pp. 211-214, doi: 10.1109/DFT.2015.7315164.
[8] Q. Lu, M. Farahani, J. Wei, A. Thomas and K. Pattabiraman, "LLFI: An Intermediate Code-Level Fault Injection Tool for Hardware Faults," 2015 IEEE International Conference on Software Quality, Reliability and Security, Vancouver, BC, Canada, 2015, pp. 11-16, doi: 10.1109/QRS.2015.13.
[9] T. Santini, et al., “Reliability Analysis of Operating Systems for Embedded SoC”. In European Conference on Radiation and Its Effects on Components and Systems (RADECS), 2015.
[10] D. Mamone, A. Bosio, A. Savino, S. Hamdioui and M. Rebaudengo, "On the Analysis of Real-time Operating System Reliability in Embedded Systems," 2020 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT), Frascati, Italy, 2020, pp. 1-6, doi: 10.1109/DFT50435.2020.9250861.
[11] I. O. Loskutov et al., "Investigation of Operating System Influence on Single Event Functional Interrupts Using Fault Injection and Hardware Error Detection in ARM Microcontroller," 2021 International Siberian Conference on Control and Communications (SIBCON), Kazan, Russia, 2021, pp. 1-4, doi: 10.1109/SIBCON50419.2021.9438916.
[12] D. Mamone, et al., "On the Analysis of Real-time Operating System Reliability in Embedded Systems," IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT), 2020.
[13] W. Mansour and R. Velazco, "SEU fault-injection in VHDL-based processors: A case study," 2012 13th Latin American Test Workshop (LATW), Quito, Ecuador, 2012, pp. 1-5, doi: 10.1109/LATW.2012.6261258.
[14] S. Azimi, et al, “Exploring the Impact of Soft Errors on the Reliability of Real-Time Embedded Operating Systems”, Electronics, 2023, DOI: 10.3390/electronics12010169.
[15] C. De Sio et al. "Analysis of Proton-induced Single Event Effect in the On-Chip Memory of Embedded Process," 2022 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT), 2022.
[16] T. J. Boger, "Rhealstone Benchmarking of FreeRTOS and the Xilinx Zynq Extensible Processing Platform," MS thesis, Dept. Elect. and Com. Eng., Temple Univ., Philadelphia, PA, USA, 2013.
[17] FreeRTOS Customisation. https://www.freertos.org/a00110.html. Accessed 2023-1-6.
[18] T. Santini, P. Rech, G. Nazar, L. Carro and F. R. Wagner, "Reducing embedded software radiation-induced failures through cache memories," 2014 19th IEEE European Test Symposium (ETS), Paderborn, Germany, 2014, pp. 1-6, doi: 10.1109/ETS.2014.6847793.
[19] M. Rebaudengo, M. S. Reorda and M. Violante, "An accurate analysis of the effects of soft errors in the instruction and data caches of a pipelined microprocessor," 2003 Design, Automation and Test in Europe Conference and Exhibition, Munich, Germany, 2003, pp. 602-607, doi: 10.1109/DATE.2003.1253674.
[20] R. Leveugle, A. Calvez, P. Maistri and P. Vanhauwaert, "Statistical fault injection: Quantified error and confidence," 2009 Design, Automation & Test in Europe Conference & Exhibition, Nice, France, 2009, pp. 502-506, doi: 10.1109/DATE.2009.5090716.