Anomaly Detection for Dependable RISC-V-based Systems in Safety-Critical Applications

Niccolò Frascarelli, Nicasio Canino, Pierpaolo Dini, Daniele Rossi, and Sergio Saponara
Department of Information Engineering, University of Pisa, Pisa, Italy

***Abstract*—Safety-critical applications require robust mechanisms to ensure dependability, thus guaranteeing safe and secure operations. Automotive and space systems operate under extreme conditions, where failure can lead to catastrophic outcomes. This study focuses on the automotive domain, aligning with the EPI project's focus on RISC-V-based embedded platforms. We propose an Anomaly Detection System (ADS) that monitors CAN bus traffic to distinguish between nominal and anomalous behavior. Our approach leverages multiclass classification algorithms, demonstrating promising memory efficiency and real-time performance. The experimental results affirm the viability of deploying lightweight ADS solutions in fault-prone automotive applications. Overall, promising performance.**

I. Topic Discussion

Dependability is the cornerstone for safe and secure systems deployed in harsh environments (e.g., extreme temperature, cosmic radiation, and electromagnetic interference), such as the space and automotive domains. Both domains require stringent performance under stress and the ability to cope with unexpected failures. Dependability encompasses both the reliability of the system and the resilience of security measures against intentional threats or unintentional faults [1].

As a case study, we selected the automotive domain, a target application also in the EPI project, driven by trends like autonomous driving and V2X (Vehicle-to-Everything) infrastructures. Using RISC-V-based embedded platforms, our work aims to address challenges such as maintaining real-time performance, reliability, and security in cost- and energy-sensitive environments. Indeed, we highlight the need for robust Anomaly Detection Systems (ADS) to enhance safety and security in automotive applications, by addressing issues such as hardware faults and cyberattacks.

II. Methodology

Our approach relies on the design of an Anomaly Detection System (ADS) integrated into a RISC-V-based embedded platform, as depicted in Figure 1. The ADS is intended to continuously monitor the vehicle Controller Area Network (CAN) bus, which is the *Target System* of our case study. In fact, most critical in-vehicle communications rely on the CAN protocol. From the data collected, a nominal *System Model* is determined offline (not while the system is running). It will be then used by the *Anomaly Detector* to classify the state of the monitored system between nominal or anomalous.

To determine the best Machine Learning (ML) model for the anomaly detector, we trained three multiclass classifiers: DTree, Support Vector Machine (SVM), and Multi-Layer Perceptron (MLP). The training process has involved feature extraction from CAN messages and rigorous cross-validation to fine-tune the model parameters. Particular emphasis was placed on optimizing the classifiers for limited memory footprints and rapid classification times.

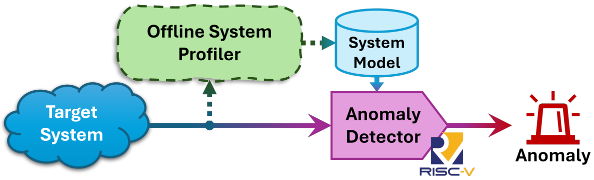

Fig. 1: *Block diagram of a typical Anomaly Detection System.*

III. Results

The evaluation of the best classifier to implement in the ADS, in addition to typical performance indexes (e.g., precision, recall, F1-score), focused on three main criteria: overall accuracy, memory footprint, and classification time. The performance of the multiclass classifiers is summarized below.

TABLE I: *Estimated resource usage and processing times.*

Classifier	Accuracy	Memory Footprint [KB]	Classification Time [ms]
DTree	0.9962	12	9.05
SVM	0.9486	240	23.45
MLP	0.9428	10	33.45

Our experiments indicate that all three classifiers are capable of operating within the tight constraints of embedded systems, typical of the automotive domain. In addition, our ADS effectively distinguishes between nominal CAN traffic and various forms of anomalous activity, thereby strengthening the potential of the system to improve vehicular dependability.

IV. Conclusion

Our study demonstrates that lightweight ADSs are feasible for RISC-V-based automotive systems. The anomaly detection system, monitoring the CAN bus, effectively distinguishes nominal from anomalous traffic under tight memory and real-time constraints. Evaluations using Decision Tree, SVM, and MLP classifiers show promising accuracy and resource efficiency. In particular, the Decision Tree classifier achieves 99.62% accuracy with a moderate memory footprint and minimal classification time, making it ideal for embedded deployment. These results confirm that robust ADS can be integrated into automotive systems without compromising dependability. In summary, this research lays the foundations for the development of more resilient automotive platforms.

References

[1] R. K. Kaur *et al.*, "Dependability analysis of safety critical systems: Issues and challenges," *Annals of Nuclear Energy*, vol. 120, pp. 127–154, 2018.

This research was partially supported by EuroHPC Joint Undertaking projects EPI-SGA2 (European Processor Initiative) and DARE (Digital Autonomy with RISC-V in Europe).