Silent Data Corruption in Robot Operating System: A Case for End-to-End System-Level Fault Analysis Using Autonomous UAVs

Yu-Shun Hsiao, Zishen Wan, *Student Member, IEEE*, Tianyu Jia, *Member, IEEE*, Radhika Ghosal, Abdulrahman Mahmoud, *Member, IEEE*, Arijit Raychowdhury, *Fellow, IEEE*, David Brooks, *Fellow, IEEE*, Gu-Yeon Wei, *Senior Member, IEEE*, and Vijay Janapa Reddi, *Member, IEEE*

***Abstract*—Safety and resiliency are essential components of autonomous vehicles. In this research, we introduce ROSFI, the first robot operating system (ROS) resilience analysis methodology, to assess the effect of silent data corruption (SDC) on mission metrics. We use unmanned aerial vehicles (UAVs) as a case study to demonstrate that system-level parameters, such as flight time and success rate, are necessary for accurately measuring system resilience. We demonstrate that downstream ROS tasks such as planning and control are more susceptible to SDCs than the visual perception stage in the perception–planning–control (PPC) compute pipeline. This observation only becomes apparent when we consider the complete end-to-end system-level pipeline, as opposed to isolated compute kernels, as previous work does. To enhance the safety and robustness of robot systems bound by size, weight, and power (SWaP), we offer two low-overhead anomaly-based SDC detection and recovery algorithms based on Gaussian statistical models and autoencoder neural networks. Our anomaly error protection techniques are validated in numerous simulated environments. We demonstrate that the autoencoder-based technique can recover up to all failure cases in our studied scenarios with a computational overhead of no more than 0.0062%. Finally, our open-source methodology can be utilized to comprehensively test the robustness of other ROS-based applications. It is available for public download at https://github.com/harvard-edge/MAVBench/tree/mavfi.**

***Index Terms*—Anomaly detection, resilience, robot operating system (ROS), silent data corruption (SDC), unmanned aerial vehicle (UAV).**

Manuscript received 14 May 2023; revised 11 September 2023 and 19 October 2023; accepted 22 October 2023. Date of publication 13 November 2023; date of current version 21 March 2024. This work was supported in part by ADA and CoCoSys, two of six centers in JUMP, a Semiconductor Research Corporation (SRC) Program sponsored by DARPA. This article was recommended by Associate Editor A. Aminifar. *(Yu-Shun Hsiao and Zishen Wan contributed equally to this work.) (Corresponding authors: Yu-Shun Hsiao; Zishen Wan.)*

Yu-Shun Hsiao, Radhika Ghosal, Abdulrahman Mahmoud, David Brooks, Gu-Yeon Wei, and Vijay Janapa Reddi are with the School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138 USA (e-mail: yushun_hsiao@g.harvard.edu).

Zishen Wan with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: zishenwan@gatech.edu).

Tianyu Jia was with the School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138 USA. He is now with the School of Integrated Circuits, Peking University, Beijing 100871, China.

Arijit Raychowdhury is with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA.

Digital Object Identifier 10.1109/TCAD.2023.3332293

I. Introduction

SILENT data corruption (SDC) has become an important problem for computing [1]. It has shown a significant threat in server scale systems [2], [3]. However, there are emerging application areas where SDCs' effects extend beyond just computational reliability into safety. Such an emerging area is autonomous vehicles where safety and reliability are critical.

Prior works have studied SDCs in the context of autonomous cars [4], [5]. However, prior work has yet to carefully examine the system-level effects of the middleware that orchestrates the entire perception, planning, and control (PPC) flow, where resiliency can be baked in to ensure SDC detection and recovery. To this end, we focus on the system-level implications of fault injection (FI) on the robot operating system (ROS) using unmanned aerial vehicles (UAVs) as a proof of concept vessel, as UAVs are agile and highly sensitive to real-time input. UAVs are predicted to have a significant market shortly due to their diversity in applications and uses [6], [7]. Nevertheless, practical safety considerations, such as performing unmanned tasks safely and without collision, impede the wide adoption of these safety-critical applications in many real-world scenarios. SDCs caused by external radiation [8] and voltage noise [9] in the computational element like the computing subsystem present a major threat to the safe deployment of UAVs [10], [11].

There are multiple error mitigation techniques, including dynamic verification [12] and redundancy [13] at the hardware or software level to improve AVs' resilience. Although current methods prove their effectiveness, they face impracticality when applied to size, weight, and power (SWaP)-constrained AVs like UAVs, primarily due to the constraints imposed by power requirements and the physical dimensions of UAV systems. Recent software techniques [14] for the resilience of convolutional neural networks (CNNs) on GPU do not apply to UAVs that typically do not have access to power-hungry GPUs onboard. Moreover, UAVs operate under stringent constraints, including limited onboard battery capacity, which imposes strict limitations on the total flight duration. Therefore, UAVs need a lightweight fault mitigation technique to prevent SDC from detouring or even crashing the UAV without compromising flight performance and availability. To this end, we set out to answer three fundamental questions.

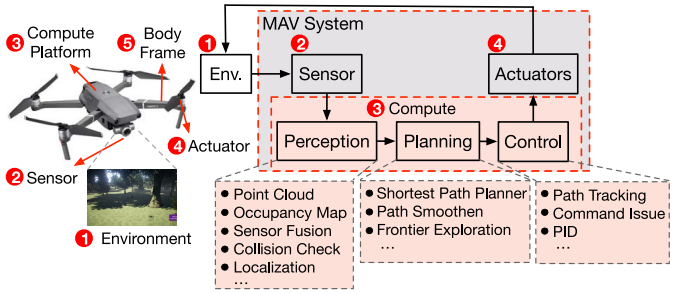

Fig. 1. End-to-end PPC computing paradigm. Each PPC stage contains multiple kernels, and we study the safety and resilience of the end-to-end pipeline.

1) What is the *SDCs' impact on system-level autonomy metrics*, such as flight time, energy consumption, and mission success rate (SR) for autonomous aerial robots such as UAVs?
2) Could conventional single, isolated-kernel SDC analysis provide similar insights as our *end-to-end fault characterization* based on system- and application-level metrics?
3) How to *enhance the safety and resilience of an autonomous robotic system against SDCs* with a lightweight mitigation technique under SWaP constraints embedded inside ROS?

To answer the first and the second question, we propose system-level metrics for evaluation and perform extensive fault characterizations (Section IV) on a real physical ROS-based autonomous system. The autonomous UAV compute consists of an end-to-end PPC pipeline (Fig. 1) that generates flight commands based on the environment in real time. The PPC pipeline is the decision-making center for a UAV to maneuver safely. An SDC could cause a UAV to detour or even crash. To analyze the impact of SDCs, we adopt the bit-flip model for FIs into the UAV's PPC pipeline and obtain quality-of-flight (QoF) metrics to quantify the impact of the faults on safety at the end-to-end whole application level.

Our findings show that application-aware metrics are essential for the resilience analysis of robotic applications. Analysis focusing on an individual computing stage without considering inter-kernel interactions leads to suboptimal insights and misguided conclusions. Prior works [15], [16], [17], [18] rely on the SDC rate to determine the vulnerability of a single compute kernel. However, a high SDC rate at a kernel-level may have a negligible impact on the QoF metrics.

For the third question, we are interested in improving UAV's safety and resilience with a lightweight mitigation technique. To this end, we propose two software-directed and lightweight enhancements for the resilience of UAV systems (Section V). Because agile robots like UAVs are constrained by SWaP, lightweight solutions are necessary. We perform data preprocessing to extract UAV's kinetics by calculating the delta value of the inter-kernel states. Based on the delta values, we perform two anomaly detection techniques. First, we perform a Gaussian-based anomaly detection (GAD) and recovery mechanism (Section V-C). This technique features a Gaussian-based range detector to exclude outliers. Second, we use an autoencoder-based anomaly detection (AAD) technique for improved UAV resilience (Section V-D). AAD adopts a neural network-based autoencoder to learn normal UAVs' kinematics and detect anomalies according to the reconstruction error of the input delta values. We show that our application-aware error detection and recovery techniques save energy by up to 1.91× than traditional redundancy-based hardware solutions [e.g., dual modular redundancy (DMR) and triple modular redundancy (TMR)] that increase the weight and form factor of UAV and lead to performance overheads.

We evaluate the effectiveness of the two detection and recovery techniques across four vastly different types of environments on two computing platforms. Our experimental results demonstrate that the Gaussian-based technique recovers up to 89.6% of failure cases, and the autoencoder-based can recover all failures in the best-case scenario. Regarding QoF metrics, the Gaussian-based technique can recover the SDC-degraded flight time by up to 63.5% and 73.0% for the autoencoder-based technique. Furthermore, our measured overhead is less than 0.0062%. Moreover, our end-to-end fault analysis framework is more generally applicable to other types of (U)AVs.

In summary, the contributions of this work are as follows.

1) We present an *end-to-end ROS-based application-aware resilience analysis framework ROSFI* to analyze robot applications' fault tolerance characteristics with proper metrics. ROSFI is seamlessly integrated with the ROS ecosystem and can be adapted for various ROS-based applications.
2) We conduct *fault tolerance characterizations of the PPC pipeline* from both kernel-level and system-level. We show that application-aware metrics are essential to understanding kernel vulnerability and fault's impact compared to the conventional isolated analysis.
3) We present *two low-cost anomaly error detection and recovery schemes* and evaluate them on different UAV configurations. By integrating anomaly error detection and recovery in ROS, We demonstrate that SDC impact on safety can be rectified in real time with negligible overhead.

II. Background and Motivation

Safety Standards: Many efforts have been dedicated to autonomous vehicle safety [19], [20]. The safety standard ISO 26262 [21] has been developed to provide guidance and safety requirements for vehicle and their systems. There are also online safety protection hardware systems developed for vehicles, such as the NXP FS4500 system for functional safety measurement. A variety of fault tolerance analyses has been performed for the computing system [4], [22] of AVs. Unfortunately, to date, there are no comprehensive standards for *autonomous* UAV assessment. Prior works mainly focus on evaluating learning-based navigation system [23], [24], [25]. However, the PPC compute paradigm-based UAV system is widely adopted in UAV systems nowadays, and its fault tolerance has not been adequately explored. Therefore, we take the first step to explore how SDCs propagate through the PPC pipeline and impact the safety of UAV systems. In this work, we define UAV reliability as the fault tolerance of autonomy kernels and UAV safety as the QoF at the application level for mission execution.

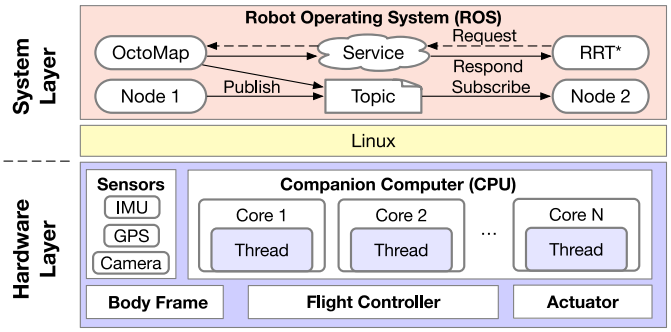

Fig. 2. System stack for a UAV. UAVs are complex cyber–physical systems with strong interdependencies between the computing and physical components. We focus on how faults in the companion computer affect the rest of the system.

System Layer: To understand how to address safety and resilience in UAVs, we must understand their complex system configuration. To this end, Fig. 2 presents the software–hardware stack of a UAV system. The system layer includes both ROS and Linux. ROS is the commonly used "operating system" to provide communication functions and resource allocation for robotic applications. Despite its name, ROS is not an operating system but a collection of robotics middleware and tools aimed at managing cyber–physical systems by providing services for heterogeneous computing, low-level device control logic, and message-passing between processes. ROS consists of multiple ROS nodes, ROS services, and a ROS master to support the functionalities and communications [26], [27]. Underneath ROS, the Linux system maps workloads to compute units and schedules tasks at runtime. Each ROS node is treated as a process that is scheduled to a thread on CPU cores.

Hardware Layer: This layer consists of sensors, a companion computer, and a flight controller. The companion computer is used to execute the PPC kernels. These kernels usually act as ROS nodes and run on a general-purpose processor (e.g., CPU). Unlike autonomous vehicles, UAVs are limited in computing resources and energy budget, and thus, it is less common to equip UAVs with power-intensive GPU. The companion computer would generate high-level flight commands (e.g., velocity in x, y, and z directions) in response to the sensor readings. The flight controller converts the high-level flight commands to low-level actuation commands to control and stabilize the UAV. In this work, we consider the faults in the companion computer and not the flight controller. The former determines the flight commands based on the real-time sensor readings, while the latter executes the commands. For instance, a corrupted yaw rotation generated by the companion computer could direct the UAV to point toward an obstacle and cause collisions. Meanwhile, the flight controller only executes the given commands without knowing the world models.

Algorithms: There has been significant advancement in perception, localization, mapping, and deep learning. Among all autonomy paradigms, the PPC computational pipeline is a widely used system [28], [29], [30]. In the PPC pipeline, the perception stage takes the sensor data and creates 3-D models to provide a volumetric representation of space, such as a point cloud [31] and occupancy map [32]. The 3-D models are then fed into the planning stage to determine a collision-free trajectory by running a motion planner [33]. Finally, based on the UAV's dynamics, the control stage follows the planned path through controllers [34].

TABLE I
COMPARISON OF FI TECHNIQUES

Abstraction Layer	Platform	Perf. (cycles/sec)	Exec. Time (1 run)	Exec. Time (1000 runs)
RTL	IVM Alpha processor RTL simulation [35]	6×10^2	4.2×10^5 hours	1.74×10^7 days
Micro-architecture	gem5 simulator [36]	3×10^6	83.3 hours	3472 days
FPGA Emulation	OpenSPARC T1 FPGA emulation [37]	1×10^7	25 hours	1040 days
Architecture	SPARC simulator [38]	6×10^7	4.17 hours	173.6 days
Software (Ours)	x86 processor [39]	3×10^9	5 mins	3.48 days

III. ROS Fault Injection

To analyze SDCs' impact on ROS, we first and foremost need an FI framework in the ROS middleware for injecting faults into the end-to-end UAV application pipeline to assess their impact systematically. This section presents ROSFI that supports FI with QoF metrics for evaluation.

A. Fault Injection Method Choices

Faults can be injected and simulated at different levels of the stack, ranging from low-level RTL [40] to high-level software [4], as shown in Table I. Although RTL simulation can accurately capture logic errors at the logic or gate levels, it requires an extremely long simulation time. In addition, it needs the RTL design or netlist of target processors, which is normally unavailable. On the other hand, software-level error injection has been widely utilized for system analysis with large vulnerability exploration space, showing significantly shorter simulation time and wide error cover range [41].

We adopt a software-level FI method to support system-level analysis, which aligns with previous fault tolerance studies for AVs by NVIDIA [4]. Software-level FI presents the best approach for an *end-to-end* study of the UAV pipeline; end-to-end implies the flow of data from the perception stage to the planning and control stages.

We assume that faults injected in ROSFI can corrupt the architectural states. Memory and caches are assumed to be protected with error correction code (ECC). Each injected fault is characterized by its location and the injected value [4]. The faults injected into the architectural states of processors can manifest as errors in the inputs, outputs, and internal states of application-level ROS nodes. ROSFI can directly inject errors in ROS node outputs by corrupting the corresponding variables based on hardware layer results. These variables are ultimately stored in different levels of storage hierarchies. Single- or multiple-bit faults cause corruption of variables when not masked in hardware. Hence, faults are injected into these memory units, and the application-level variables are corrupted accordingly to emulate the faults.

The need for a software-based approach is justified by Fig. 3, which demonstrates the FI execution time of single-kernel and UAV experiments at different layers. One UAV run includes hundreds of single-kernel executions for the UAV

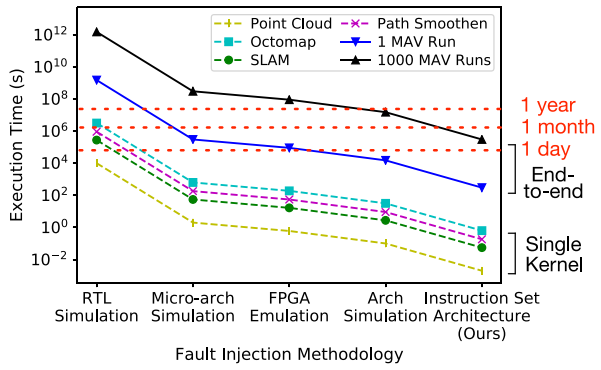

Fig. 3. Comparison of techniques at layers of abstraction. Performing end-to-end analysis requires fast execution time.

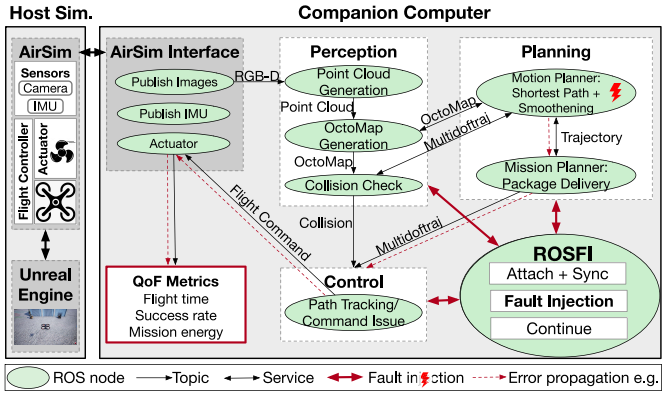

Fig. 4. Overview of the end-to-end application-aware ROSFI resilience analysis framework.

autonomous navigation experiments, which take 5 min per run or 3.48 days per scenario (i.e., 1000 runs) with our software-level execution. Consequently, it is infeasible for extensive fault analysis involving thousands of FIs at the lower levels of the abstraction layers.

B. ROS Fault Injector Implementation

Fig. 4 illustrates the FI infrastructure of a ROS-based UAV system, including environment and UAV simulation on the host simulator and the UAV's PPC pipeline integrated with ROSFI on the companion computer. Each PPC stage contains one or multiple ROS nodes, and each ROS node comprises a single compute kernel, such as point cloud generation or motion planner. ROS node communicates through ROS topics (one-to-many communication) and ROS services (one-to-one communication). The ROSFI is built as a ROS node to maintain our framework's portability, which leverages the ROS communication protocol and Linux system call.

To establish UAV experiments, we leveraged an open-source ROS-based UAV simulator, MAVBench [28]. MAVBench includes unreal engine (UE) to simulate the surrounding environment, AirSim simulator [42] to capture a UAV's dynamics and kinematics, and PPC computational pipeline to generate flight commands in real time. The AirSim interface allows the PPC pipeline to access the sensor data and send back the flight commands to the flight controller in the AirSim simulator. The PPC pipeline processes the sensor data and generates flight commands continuously until the mission is complete. Finally, the mission QoF metrics are recorded. Although we use UAVs as an example in our framework, the fault analysis methodology is broadly applicable to any ROS-based use case.

Fig. 4 also illustrates an error propagation example within the system. For instance, when a fault is injected at the *Motion Planner* kernel and manifests as a corruption of execution results (i.e., *Multidoftraj* and *Trajectory*), which eventually corrupts a flight command and impacts the overall QoF. The framework works on x86/Linux platforms. Fig. 5 shows the instruction-level FI details of ROSFI. Each oval node is a ROS node. The figure illustrates the FI sequence using an example for ROS node 2. During the system initialization phase, the ROSFI node publishes its process ID (pID) to all the other nodes and subscribes to their pID. Thus, the ROSFI

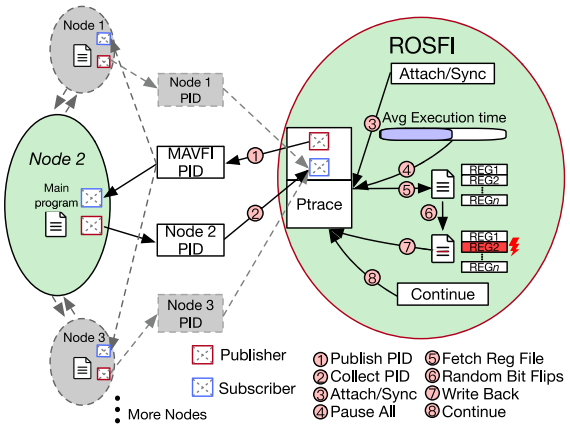

Fig. 5. Design details of the ROSFI FI node and its interactions with other ROS nodes.

node can attach and manipulate the other ROS nodes in the system via the *ptrace* system call supported by Linux. The *ptrace* system call allows synchronization and manipulation of processes' register files with much less overhead than the ROS communication protocol.

ROSFI is the first FI framework built on top of *ptrace* system call and ROS. It emulates SDCs that occur in the processor's functional units (e.g., arithmetic and logic units) by introducing transient bit-flips at the source or destination register of only one dynamic instruction, which is known as *instruction-level FI* [43], [44], [45]. We do not consider faults in the memories or caches as they can be protected by ECCs in safety-critical applications. ECC is used to protect memory for robots that use TX2-level hardware (as considered in this article). We also assume no faults in the processor's control logic, which constitutes only a small portion of the processor. This is in line with previous fault analyses [16], [17], [46]. Hence, while our approach does not cover all the FI sites, it provides us with quality early-stage, end-to-end insights.

To inject faults, ROSFI selects a random time point to pause all the nodes during the simulation of real-time ROS applications. All ROS nodes' execution will be stopped before FI, ensuring that every node follows the original executive order. After all the nodes have stopped, the general-purpose and floating-point register files of the target node (i.e., node 2)

TABLE II
COMPARISON BETWEEN ROSFI AND PRIOR FI METHODS

	LLFI [44]	**PINFI** [43]	**CLEAR** [48]	**SASSIFI** [15]	**DriveFI** [4]	**ROSFI (This work)**
Autonomous Vehicle	✗	✗	✗	✗	✓(Vehicle)	✓(Drone)
Support ROS	✗	✗	✗	✗	✗	✓
Platform	CPU	CPU	CPU	GPU	GPU	CPU
Injection level	IR-level	Instruction-level	RLT-level	Instruction-level	Instruction-level, Source-level	Instruction-level, Source-level
Single Bit-flip	✓	✓	✓	✓	✓	✓
Double Bit-flips	✓	✗	✗	✓	✓	✓
Multiple Bit-flips	✗	✗	✗	✗	✓	✓

are fetched via the *ptrace* system call, with the instruction pointer register decoded to access the current operating register. The number of registers being accessed by an instruction ranges from 0 to 2. If the value is zero, ROSFI resumes all ROS nodes' execution and repeats the above steps to obtain a new instruction. For more than one register under operation, ROSFI randomly chooses one register to inject. For the source register, according to the user-defined injection configuration, a single bit-flip or multiple bit-flips are introduced. For the destination register, before FI, ROSFI would step toward the next instruction to allow the current instruction to finish the write, which avoids the corrupted destination register being overwritten by the current instruction. After FI, the corrupted register is written back to the target node's register files, and all nodes are notified to resume the execution. The faults injected into the registers of the processors could manifest as errors in the inputs, outputs, and internal states of the computational kernels. To better understand error propagation among PPC stages, ROSFI can inject errors into the inter-kernel states (the input of the other kernel) via *source-level FI* [4].

ROSFI can inject either single or multiple bit-flips simultaneously. In a previous study [47], it was shown that a single bit-flip is good enough for fault analysis since it can capture first-order vulnerability characteristics as well as multiple-bit-flips analyses. Therefore, for the analysis results in this article, we mainly focus on a single bit-flip. For simplicity and clarity, we refer to single-bit-flip FI in the rest of this article unless multiple bit-flips are specified.

ROSFI has the potential to extend to cover the memory and control logic of the processor. For memory, in this article, we assume memory and caches are protected with SECDED codes. The faults injected in instructions may result in accessing the wrong data for computation, thus corrupting the variables. Faults in memory will result in corrupted computation data as well. These variables are ultimately stored in different levels of storage hierarchies (e.g., registers or caches). For control, since ROSFI obtains the whole instruction, it is able to modify the opcode of the instruction, thus the control logic.

C. Comparison to Prior Art

SDCs and resilience analysis have been studied for single kernels on CPU and GPU, as shown in Table II. However, prior methods [15], [16], [17], [18] focus on the SDC rate of a single kernel, which does not directly translate to the impact of SDCs on UAVs' QoF metrics. On the one hand, more recently, DriveFI [4] explored the resilience impact of SDCs for autonomous driving systems on power-hungry GPU platforms.

However, there currently does not exist an FI framework to analyze the resilience of ROS-based applications where the ROS nodes typically run on CPU with ROS [49]. Furthermore, a difference between prior autonomous vehicles resilience analysis studies and UAVs is that the *compute* environment of a UAV diverges from the recent trend in autonomous vehicles. While we find that most recent systems and tools are migrating to GPUs for increased parallel processing and DNN acceleration (and thus, may require GPU-centric resilience analysis tools, such as SASSIFI [15] or NVBitFI [50]), UAVs continued to operate in the CPU realm for flexibility and lower energy needs. Currently, UAVs are intrinsically associated with CPUs due to the reliance on ROS (Section II).

IV. END-TO-END PPC PIPELINE FAULT TOLERANCE

This section presents the fault tolerance analysis at two granularity levels, i.e., single-kernel level and application-aware system-level performance. We explore how errors would impact a single kernel and propagate through the whole PPC pipeline to affect UAV QoF metrics. Through the comparison, we observe that isolated, single-kernel analysis (as is common practice) provides different or even opposite insights on the vulnerabilities of kernels than the application-aware analysis, which shows that end-to-end application-aware fault analysis is crucial to capturing SDCs' impact at the system level.

Metrics: For single-kernel level analysis (Section IV-A), we use benign, crash, hang, and SDC rates to evaluate the resilience of representative autonomy kernels. For application-level analysis (Section IV-B), we use UAV QoF (i.e., mission SR, time, and energy) to evaluate end-to-end system performance and resilience characteristics.

A. Kernel-Level Fault Tolerance Analysis

To prove the importance of application-aware fault analysis, we first conduct the single-kernel analysis with instruction-level FI. The single-kernel FI flow is similar to prior FI works [51]. We show that conducting similar analyses in a complex PPC pipeline can lead to misguided conclusions, specifically for UAV systems.

We evaluate the common kernels in the PPC pipeline, including *OctoMap* [52] for the perception stage, three sampling-based motion planners [53] (i.e., *RRT*, *RRTConnect*, and *RRT**) for the planning stage, and PID controller [28]

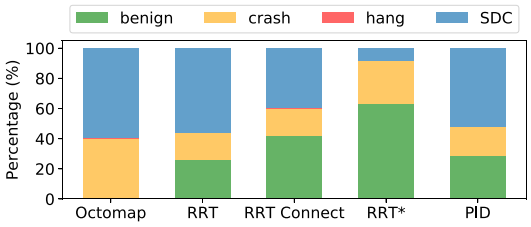

Fig. 6. Conventional isolated single-kernel analysis.

for the control stage. For each kernel, we perform instruction-level FI for 5000 runs in total. Each kernel is run without FI to obtain the error-free golden results. With FI, there are four outcomes: 1) execution results same as the golden results (i.e., *benign*); 2) execution exceptions (i.e., *crash*); 3) infinite execution time (i.e., *hang*); and 4) execution results different from the golden results (i.e., SDC) [54].

From a single-kernel perspective, the perception stage is the most critical when an SDC manifests. Fig. 6 shows that most compute kernels are more than 25% *benign* error-tolerant except for *OctoMap* at the perception stage. This is because SDC could easily manifest at the output (*Octree*) with noisy values for the *OctoMap* kernel. Hence, *OctoMap* is less resilient than sampling-based planning and PID control algorithms. On the other hand, the path planning kernels (*RRT*, *RRTConnect*, and *RRT**) are all sampling-based algorithms, which are known for their high efficiency and performance for low-dimensional planning. Injected faults should not affect output results as long as the corrupted waypoint is not sampled. The more waypoints we sample, the higher the probability the planning algorithms could sample a corrupted waypoint, resulting in SDC. *RRTConnect* runs two *RRT* algorithms from both start and goal positions, ending up with fewer sampled waypoints than *RRT*. *RRT** is the optimized version of *RRT* algorithm to find the shortest path by selecting even fewer waypoints than *RRT* and *RRTConnect*, making *RRT** having the least SDC. The PID algorithm at the control stage also experiences around 25% *benign* cases as the PID has a simple self-healing mechanism to clip data outside of a bounded range.

B. End-to-End, System-Level Fault Tolerance Analysis

Compared to the single-kernel fault analysis, we next conduct application-aware fault analysis based on our ROSFI framework. This end-to-end system-level characterization investigates how kernel errors would propagate through PPC pipelines and impact UAV performance. We assess the performance of the UAV using QoF metrics, encompassing key aspects such as flight SR, distance, time, and energy. Ultimately, these metrics matter from an "application" or system-level perspective. The SR quantifies the ratio of successful missions to the total number of flight runs. We define a mission as successful when the UAV reaches its destination without encountering any collisions. Conversely, a failure occurs when the UAV either collides with obstacles or is unable to identify a viable path to its intended destination. Flight time represents the total duration required for the UAV to reach its designated destination. Similarly, flight energy signifies the overall energy expended by the UAV to reach the destination. Since rotors dominate mission energy (∼95% [28]), flight energy positively correlates with flight time. It is worth mentioning that with reduced single-mission flight energy (E), the number of completed missions (N) under a battery charge (E_b) and SR will increase through $N = SR \times (E_b/E)$.

We adopt the instruction-level fault injector supported by ROSFI to corrupt the PPC kernels. In our default settings, the PPC pipeline includes *Point cloud generation, OctoMap, Collision check* for perception, *RRT** for planning, and *PID* for control. Two other common planning algorithms are evaluated at the planning stage, i.e., *RRT* and *RRTConnect*. Each kernel has experimented with 100 FI runs. Besides FI, 100 error-free experiment runs are defined as *Golden*. In each experiment, all kernels in the PPC pipeline would be launched by ROS to complete a given navigation task. Only one of the kernels would have a one-time FI during each flight mission for FI runs. We achieved a 4.45% error margin with a 95% confidence level with 100 experiment runs per configuration. Without loss of generality, we limit our discussion to a navigation task in the *Sparse* environment here. More results are demonstrated in Section VI.

Counter to the single-kernel perspective, from an end-to-end application perspective, the perception stage is the least critical when an SDC manifests. Prior works generally tend to focus on error resilience of the perception stage [1], [14], [55]. These are aligned with the single-kernel analysis, which shows that *OctoMap*, the main algorithm for perception, has the highest percentage of SDC among the evaluated kernels. However, as we show, for the perception stage both *Point Cloud Generation* and *OctoMap* have little to negligible impact on the system metrics, as shown in Fig. 7.

The reason why *OctoMap* has the least impact on QoF metrics is that even if an occupied voxel is corrupted and mistaken as a free voxel, the presence of all other surrounding voxels as occupied ensures that the UAV can still accurately determine the locations of obstacles. This holds as long as the resolution of the *OctoMap* is adequate, allowing the UAV to make the correct decisions regarding its flight path. This counter-intuitive insight underscores the significance of conducting comprehensive end-to-end analysis.*Collision Check* is the critical kernel in the perception stage since a false alarm can lead to trajectory replanning or collisions.

From the end-to-end application-level perspective, planning, and control are more critical than perception, counter-intuitive to the single-kernel analysis. While the SDC percentages of planning and control kernels are lower than *OctoMap*, corrupted outputs (e.g., yaw, roll, pitch, and velocity) from these two stages can directly lead to a detour or crash of the UAV. From Fig. 7(a), even though the average flight time is similar, the range of *RRT, RRTConnect, RRT*, and PID* is much wider than *Octomap* and *Golden*. The error propagation of the corrupted execution results could greatly increase the flight time by up to 57.3% and even lead to degradation of SR by up to 8% as shown in Fig. 7(c). Hence, the planning and control stages are more critical than the perception stage from an end-to-end application perspective.

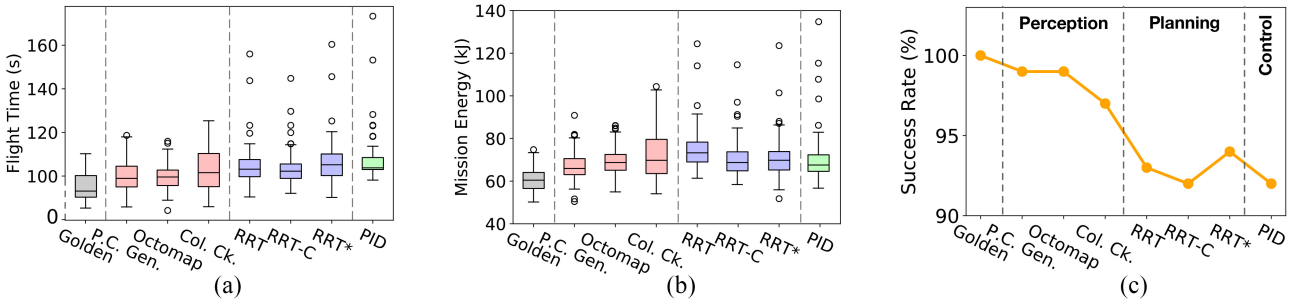

Fig. 7. Application-aware system-level end-to-end resilience analysis (flight time, energy, and SR) with ROSFI framework. (a) Flight time. (b) Flight energy. (c) Flight SR.

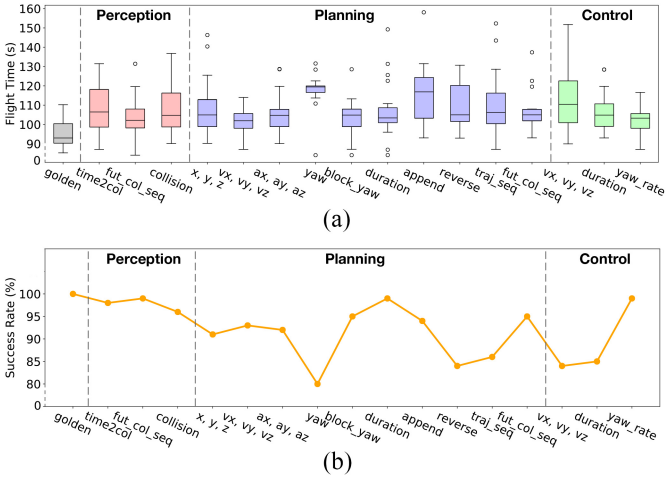

Fig. 8. End-to-end fault tolerance analysis. (a) Flight time. (b) Flight SR.

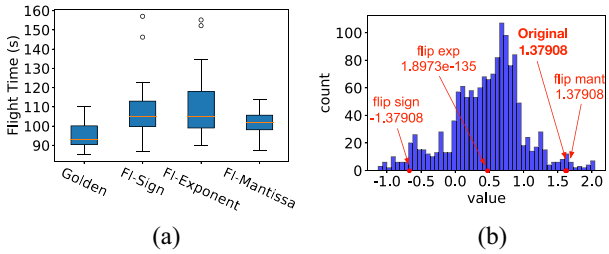

Fig. 9. Impact of FI at different data fields. (a) Flight time. (b) Bit-flip analysis.

TABLE III
QOF METRICS WITH A DIFFERENT NUMBER OF BIT-FLIP INJECTIONS

Bit Flips	Flight Time (s)	Flight Distance (m)	Total Energy (kJ)	Number of Re-plans	Success Rate (%)
0 (golden)	94.6	49.5	51.4	3.71	100
1	105.6	55.9	57.2	4.07	92
3	107.3	56.8	58.7	4.19	91
5	111.8	59.4	61.1	4.28	89

C. Error Propagation Across PPC Stages

To understand error propagation across kernels, we analyze the impact of corrupted inter-kernel states in the PPC pipeline. This provides insights to improve the PPC kernels and facilitate error detection and mitigation in Section V. We do source-level FI for this inter-kernel error propagation study with 100 navigation task runs for each evaluation.

As shown in Fig. 8, inter-kernel states exhibit different resilience to faults and have diverse impacts on UAV QoF metrics based on their functionality. For example, in the perception stage, *future_collision_seq* is much more robust than *time_to_collision*, whose QoF metrics noticeably vary when compared to the golden run. Faults in *time_to_collision* can skew the UAV's perceived distance to obstacles. Similarly, data corruption of (x, y, z) and *yaw* of waypoints planned by motion planner can lead to a wrong direction or crash into obstacles, and faults in (vx, vy, vz) could make the UAV fail to keep track of a trajectory. As a result, the distorted trajectory leads to collision or increased flight time and mission energy.

Bit-flips in different data fields have distinct levels of impact on UAV performance. Prior works have evaluated data field impact on the processor and neural network [1], and we further corroborate this in end-to-end UAV systems from the application-level perspective. We conduct source-level FI at the float64 inter-kernel states (x, y, z), which contains 1 sign bit, 11 exponent bits, and 52 mantissa bits. Faults in sign and exponent fields have a greater impact on the UAV's resilience and result in increased flight time, energy, and failure cases, as shown in Fig. 9(a). Faults in the sign and exponent will result in a greater change in the inter-kernel states than faults in the mantissa field. For example, a single bit-flip at the exponent and sign could corrupt 1.38 to 0 and -1.38, respectively, as illustrated in Fig. 9(b). The huge differences show that sign and exponent fields are more critical to the UAV system when an SDC manifests and propagates through the PPC pipeline. We further leverage this insight in lightweight UAV anomaly detection in Section V.

To compare single bit-flip with multiple bit-flips, we evaluate the performance impact with multiple bit-flips FI as shown in Table III. In this experiment, 100 FIs are performed for 1, 3, and 5 bits, respectively, at *(ax, ay, az)*, which are the output states of the planning stage. From 1-bit to 5-bits FI, the average flight time and energy slightly increase by 6.2 s and 3.9 kJ, respectively, and the SR decreases by 3%. Since more bit-flips are more likely to affect the sign and exponent fields, the value changes could be more dramatic with multiple bit-flip FI. However, the slight differences between single bit-flip and multiple bit-flips also show that single bit-flip can capture the first-order vulnerability characteristics as shown in the prior work [47].

V. Error Detection and Recovery

To enhance safety and resiliency, we further explore the detection and recovery technique based on the observations from ROSFI. As the conventional redundancy-based hardware protection introduces significant overhead, we propose two software-level low-overhead anomaly detection and recovery schemes for UAVs. The proposed schemes detect anomalous behavior of the inter-kernel states in the PPC pipeline and cease the error propagation, ensuring UAV's safety.

A. Overview of Detection and Recovery

Due to the low overhead and high effectiveness of anomaly detection, it has been used to distinguish anomaly from normal data distribution in many applications [56]. However, there is no effective general anomaly detection technique for different domains. Moreover, autonomous machines are complex systems that typically involve multiple kernels' heterogeneous computing. It is infeasible to separate normal data from anomaly based on the system's input (e.g., sensor readings) and output (e.g., flight commands). The heterogeneity also makes it hard to extract information from the system for anomaly detection. As a consequence, no prior work has focused on anomaly detection to enhance the resilience of autonomous aerial vehicles.

We propose two anomaly detection techniques to detect SDC that could cause a safety hazard for UAVs, including Gaussian- and autoencoder-based techniques. It is observed that both techniques can greatly enhance the safety and resilience of UAVs with a low computational overhead. Fig. 10(a) shows the proposed anomaly detection and recovery scheme for UAVs. According to the analysis in Section IV-C, the inter-kernel states as shown in Fig. 8 are monitored for the anomalous SDC. The monitored states pass through a data preprocessing module to increase the detection performance while further reducing the computational overhead. After data preprocessing, the processed states go into either of the proposed anomaly detection techniques for supervision.

Error recovery is a feedback loop from the detection modules to the PPC pipeline. Once an anomalous behavior is detected, an alarm signal will be raised by the detection modules, triggering the recomputation of the corresponding stage, which prevents the corrupted inter-kernel states from propagating to the other kernels. The proposed detection and recovery system can greatly increase the resilience of UAV's PPC pipeline against SDCs that degrade the safety and flight performance of UAVs. Our approach focuses on SDC as ROS node *crash* can be detected by the ROS system. The ROS master node would restart the node automatically if it crashes [57].

B. Data Preprocessing

In Fig. 10(a), the monitored inter-kernel states from the PPC pipeline are processed in the data preprocessing block before being sent to the anomaly detection block. Data preprocessing has two steps, including data format transformation and delta calculation. First, for data format transformation, the sign and exponent bits of *float64* states are transformed into 16-bits integer states. Since SDC at the mantissa bits of *float64* are insignificant as shown in Section IV-C, only the sign and exponent bits are monitored to reduce the detection overhead. Second, the deltas of the incoming states are calculated. We define delta as the number of value changes from the previous time point to the current time point for an inter-kernel state.

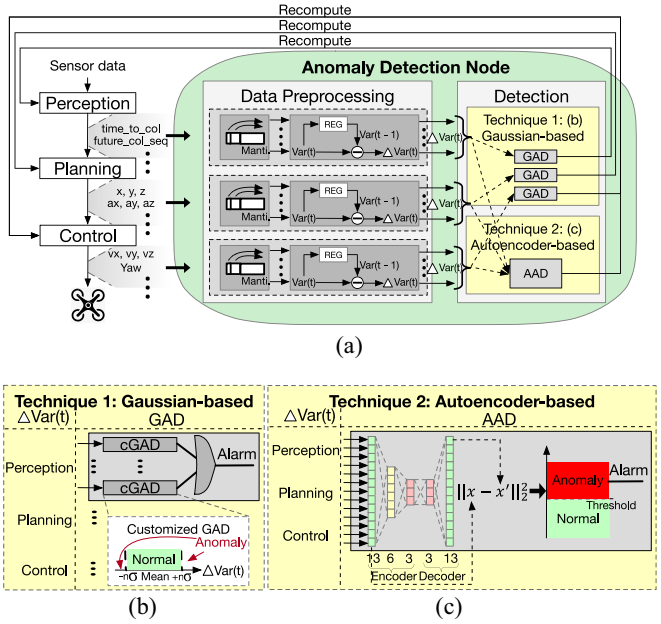

Fig. 10. Proposed anomaly detection and recovery scheme for UAV computational pipeline. (a) Overview. (b) Gaussian-based. (c) Autoencoder-based.

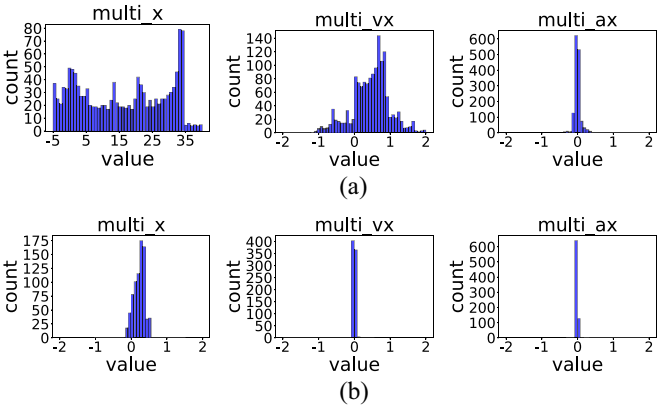

Fig. 11. Histogram comparison between the states' value and delta after data processing. (a) Histogram of the states' value. (b) Histogram of the states' delta.

Fig. 11 shows the insight of using the state's delta for anomaly detection. For most states, the value could have either uniform or Gaussian distribution. However, the uniform value distribution is not well suited to GAD, leading to very low detection accuracy. The uniform distribution can be transformed into a Gaussian distribution by calculating the states' delta, leveraging the inter-kernel states' temporal dependency. Furthermore, the state's delta range is much smaller than the original data. For instance, as shown in Fig. 11, the range of *multi_x*, *multi_vx*, and *multi_ax* states is reduced by 98%, 94%, and 76%, respectively, making the differences between normal

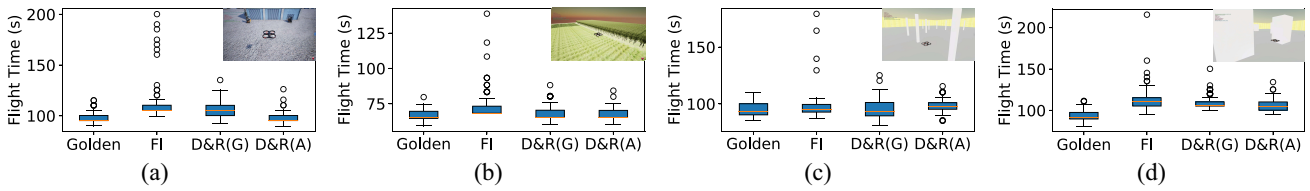

Fig. 12. Effectiveness of the proposed anomaly detection and recovery schemes in terms of flight time. D&R(G) and D&R(A) represent the Gaussian-based and autoencoder-based schemes, respectively. (a) UE Factory. (b) UE Farm. (c) Sparse. (d) Dense.

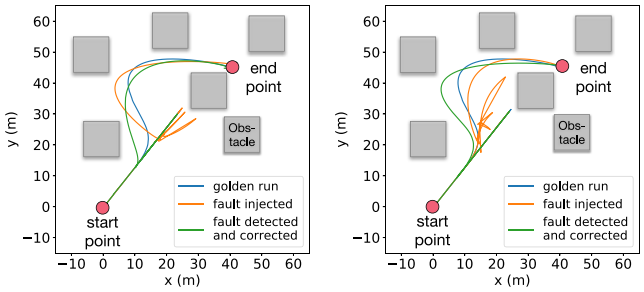

Fig. 13. Trajectories of the golden run, with FI, with both FI and error detection and recovery. (a) FI in perception. (b) FI in planning.

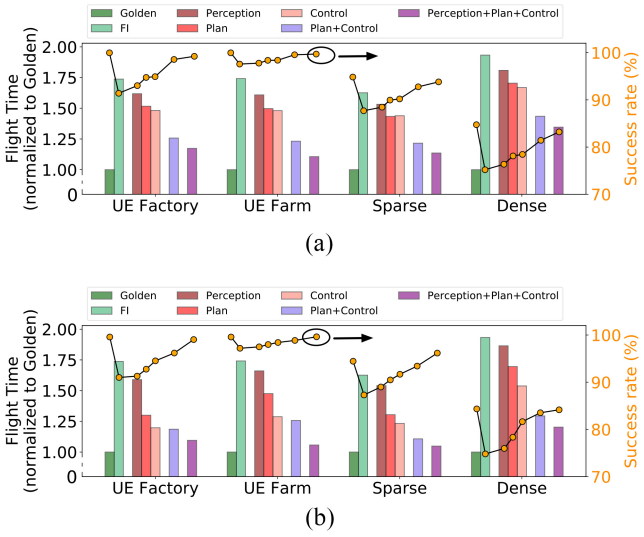

Fig. 14. Worst-case QoF metrics with different error detection and recovery stages (results normalized to golden run). (a) GAD and recovery (flight time). (b) AAD and recovery (flight time).

and anomaly data even larger. Thus, data preprocessing can increase anomaly detection performance while decreasing the overhead of detection.

C. Gaussian-Based Anomaly Detection

Fig. 10(b) shows the design of the GAD. Each PPC stage has a corresponding GAD that consists of several customized GAD (cGAD) for each inter-kernel state. If the value of an incoming state is outside the range of its normal data distribution, its cGAD will send out an alarm. The alarms from each cGAD are gathered for each PPC stage, respectively. An alarm from a GAD would trigger the recomputation path of its corresponding stage, stopping the error propagation to the next stage.

The Gaussian model parameters (i.e., mean and standard deviation) for each cGAD are estimated as follows:

$$M_k = M_{k-1} + (x_k - M_{k-1})/k \tag{1}$$

$$S_k = S_{k-1} + (x_k - M_{k-1})(x_k - M_k) \tag{2}$$

where k is the number of samples, M_k is the mean value for the k samples, and S_k is an auxiliary term used to compute standard deviation σ. At initialization, we introduce and set the terms $M_1 = x_1, S_1 = 0$. The parameters are updated online with the recurrence formulas above for new incoming data x_k [58]. For $k \geq 2$, the standard deviation σ can be derived by $\sigma = \sqrt{S_k/(k-1)}$. Whenever the value of the incoming data is n sigma away from the mean value, the alarm of the cGAD will be raised. The number of sigma n is a configurable variable that can be optimized based on the complexity of the flight task and environment. To ensure the Gaussian models have sufficient samples before starting anomaly detection, we first have the model updated with error-free training environments.

D. Autoencoder-Based Anomaly Detection

Fig. 10(c) shows the AAD. The AAD block collects the processed states from all PPC stages as input. An alarm will be raised and triggers the recomputation of the control stage if an anomaly is detected. The proposed autoencoder comprises an encoder with three fully connected layers and a decoder with two fully connected layers. The encoder has an input layer of 13 neurons, a hidden layer of 6 neurons, and an output layer of 3 neurons. The decoder has an input layer of 3 neurons, which takes the compressed data from the encoder, and an output layer of 13 neurons. The output of the decoder represents the reconstructed input data. The reconstruction error is the difference between the input and output of the autoencoder. We use the mean squared error during the unsupervised training as the reconstruction error. If the reconstruction error is beyond the threshold at the inference phase, the alarm will be raised. The threshold is the upper bound of the reconstruction error in the error-free golden run.

Rather than a separate Gaussian-based detection module for each PPC stage, we use a single autoencoder for the whole PPC pipeline to leverage the correlation among the inter-kernel states. Once an anomaly is detected, the alarm triggers the recomputation of the control stage. In this way, the autoencoder scheme achieves higher detection performance while reducing the recomputation overhead as shown in Section VII-D.

E. Recovery Scheme

Once an anomalous inter-kernel state is detected, the recomputation path will be triggered to cease the error propagation. The corresponding compute stage fetches the newest data from the previous compute stage or sensor and regenerates the results. Take the navigation task as an example. If an alarm is raised in the perception stage, the stage starts to recompute and fetch the newest RGB-D camera data. Then, *Point Cloud Generation*, *Octomap*, and *Collision Check* kernels process the data and generate results for the following stage. Similarly, if an alarm is raised in the planning stage, the planning algorithm will fetch the latest occupancy map from the perception stage and plan a new trajectory. Finally, the flight command is monitored at the control stage before being sent back to the UAV. If an alarm is raised, the control stage will abandon the current anomalous waypoint and fetch the next waypoint of the trajectory, generating correct flight commands.

F. Anomaly Detection and Recovery on ROS

The anomaly detection and recovery scheme are built as a ROS detection node. This node contains the data preprocessing and anomaly detection functions (as explained previously). The detection node subscribes to the topics containing the inter-kernel states in the PPC pipeline as input and publishes recomputation signals to the corresponding stages if the detection function raises the alarm. The detection node can thus continuously supervise inter-kernel states of the PPC pipeline, avoiding error propagation among kernels and thus increasing the resilience of UAV's computational pipeline with negligible overhead.

VI. Experimental Setup

Hardware-in-the-Loop Simulator: We used a closed-loop simulator as the experimental platform [28], including UE to simulate the scenarios and AirSim to capture the UAV's kinematics. Sensors, including RGB-D camera and IMU, used in the experiments, are common for UAVs. An Intel i9-9940X CPU and an NVIDIA GTX 2080 Ti GPU are used as the host machine to simulate environments and the UAV. The companion computer has a CPU that takes sensory data and generates flight commands for UAVs.

Training Environments: To create a training dataset for the autoencoder-based technique, we built an environment generator with configurable parameters (i.e., obstacle density and size of obstacle). The obstacle density is defined as the probability of a 10 * 10 grid spawned with an obstacle. Each obstacle is a cuboid with n * n and infinite height (n is in [1, 10]). [*obstacle density, size of obstacles*] is defined as an environment configuration pair. We collect data from randomized environments with the combinations of two obstacle densities (0.05 and 0.2) and two sizes of obstacles (3 and 5). Therefore, there are four configuration pairs in total, and each is run 25 times. A random seed is used to randomize the environment. For the Gaussian-based technique, the Gaussian models are updated with the same error-free training environments.

TABLE IV
Flight SR in Four Evaluation Environments

Environment	Factory	Farm	Sparse	Dense
Golden Run	100.0%	100.0%	95.0%	85.0%
Injection Run	91.7%	97.3%	88.3%	75.3%
Gaussian-based	98.7%	99.3%	94.3%	83.0%
Autoencoder-based	99.3%	100.0%	95.0%	84.7%

Evaluation Environments: The anomaly detection and recovery schemes are evaluated in four environments, which are unknown to the UAV. So we are not evaluating on training data. The *Factory* and *Farm* are provided by UE, representing common navigation scenarios with blocks, walls, and hedges. We generate the *Sparse* with [0.05, 3] and the *Dense* with [0.2, 5] using our environment generator. The random seed is fixed.

Overheads: The QoF metrics do not include the FI time since the ROS nodes are paused during FI. In terms of simulation time, ROSFI only takes less than 5 ms for one-time FI, which is negligible for a typical flight mission that takes more than 100 s. For anomaly detection and recovery runs, we quantify the overhead of Gaussian and autoencoder-based techniques in Section VII-D.

VII. Evaluation

To evaluate the anomaly detection and recovery scheme, we run 100 error-free simulations for each environment as the baseline (golden run). Then, we conduct 900 single-bit injections at instruction level for each environment, including 300 runs for each setting (i.e., FI, *detection and recovery with Gaussian [D&R(G)]*, and *detection and recovery with autoencoder [D&R(A)]*), as shown in Fig. 12. In each setting, we have 100 FIs for each PPC stage. Each run includes a one-time single-bit injection. A total of 1000 runs is chosen where each run takes about 5 min. Even though ROSFI introduces a negligible overhead of only 5 ms, the experiment time is a limiting factor for the total runs.

A. Safety Metrics

Improvement of SR: Table IV shows the SRs of UAV flights across four environments. In the FI runs, the SR drops 9.7% in the *Dense* environment. Faults may easily cause collisions or fail to find a collision-free path in complex environments. In contrast, *Farm* is an obstacles-free environment. Even if a UAV detours from its path, there are more feasible paths toward the destination than a complex environment. With the anomaly detection and recovery scheme, Gaussian- and autoencoder-based techniques recover up to 89.6% and 100% (fully recover) of failure cases, respectively. Generally, the autoencoder recovers more failure cases than the Gaussian-based scheme and increases the SR close to or the same as the error-free runs.

Improvement of Flight Time: Fig. 12 shows the flight time of all successful cases in Table IV across four environments. Similar to Section IV-B, the FI runs result in a much wider range of flight time than the golden run and increase the flight time by 73.8%, 74.2%, 62.6%, and 93.3% in the worst case for each environment, respectively. However, with GAD and recovery, many outliers can be recovered, and the worst-case

flight time is recovered by 56.4%, 63.5%, 49.0%, and 58.7%. On the other hand, the autoencoder-based technique recovers most of the outliers and can recover the worst-case flight time by 64.2%, 68.4%, 57.8%, and 73.0%, outperforming the Gaussian method.

Comparison of Gaussian-Based and Autoencoder-Based Schemes: The autoencoder-based technique consistently outperforms the Gaussian-based technique in SR and QoF metrics. The reason is that the autoencoder can leverage the correlation among the inter-kernel states; thus, it can detect the subtle discrepancy of the states. However, the Gaussian-based technique does not have correlation information among states. Therefore, it can only detect each variable separately, which may fail to detect anomalies if the corrupted data is still inside the range of the normal data distribution.

We provide both methods in the ROSFI framework. The Gaussian method serves as a practical and efficient solution for anomaly detection. It requires minimal data collection to update the standard deviation and mean values for each inter-kernel state. This simplicity and real-time adaptability are especially valuable in scenarios where immediate anomaly detection is critical, as it minimizes computation and overhead. The autoencoder method, while more effective in terms of detection accuracy, necessitates offline training, making it more suitable for scenarios where detection accuracy is paramount and time constraints permit offline model training. In essence, the choice to utilize both methods stems from pragmatic consideration of the diverse needs in UAV operations.

Comparison of Environments: More dense environments with a higher density of obstacles make it difficult to recover from errors. For the *Dense* environment, a UAV has more complex trajectories to follow and more dynamic actions in response to the obstacles, making the range of the variable distribution wider. The wider distribution increases the number of false-negative detections. Thus, there is still a 20.1% gap between autoencoder-based recovery results and golden for the worst case. On the other hand, for the obstacle-free *Farm* environment or *Sparse*, the autoencoder-based technique can achieve a similar performance as the golden run.

B. Trajectory Analysis

To show the impact of faults and the effectiveness of our detection and recovery schemes, we visualize UAV's trajectories in the *Dense* environment. We present the trajectories with the autoencoder-based technique, while the Gaussian-based technique has similar results when successful.

Fig. 13 shows the scenario in which a single-bit injection in the PPC stage can lead to a flight detour and how the detection and recovery scheme improves the flight. Without FI (blue curve), the UAV takes off at the start point and flies toward the endpoint in the beginning phase. Then, when facing an obstacle, it stops at a safe distance and replans a new trajectory that flies back slightly and bypasses the obstacle.

When faults corrupt critical inter-kernel states, such as the coordinate of a waypoint, the path may be distorted. The UAV may not stop until it faces an obstacle (orange curve), which causes the UAV to fly back or replan its trajectory. The more often the UAV replans and detours from its path, the longer it takes to reach the destination, which increases the flight time by 21.9% and 24.5% for Fig. 13(a) and (b), respectively. With the detection scheme, the corrupted waypoint will be abandoned once an anomaly is detected. The alarm raised by the detection module triggers the stage recomputation. Therefore, the UAV would follow a better trajectory (green curve) without detour, which recovers the QoF metrics.

C. Anomaly Detection and Recovery

To evaluate error detection effectiveness in different PPC stages, we experiment with the anomaly detection and recovery scheme for certain compute stages.

Single-Stage Detection and Recovery: We first experiment with anomaly detection and recovery by only detecting a single pipeline stage. As shown in Fig. 14, the Gaussian-based technique recovers the flight time by 16.2%, 29.9%, and 34.7% and the autoencoder-based technique recovers the flight time by 20.1%, 59.3%, and 73.2% for PPC, respectively, along with the SR improvement and flight energy savings, in *Factory* environment. A similar trend has been demonstrated in the other three environments. Both techniques show that the flight time can be recovered the most by detecting the faults that happened in the control stage. The reasons are twofold. First, the planning and control stages are more vulnerable to faults from the system perspective. Second, any error propagated from previous stages has to pass through the control stage before corrupting the flight command. The evaluation of the individual stage lines up with the analysis in Section IV-B.

Multistage Detection and Recovery: To understand how different stages affect anomaly detection and recovery, we apply the scheme to multistages, namely, the planning-and-control (PC) stage and all PPC stages. The Gaussian method recovers the flight time by 65.1% and 76.5%, and the autoencoder-based recovers by 74.8% and 87.1% for PC and PPC, respectively, along with the SR increase and fight energy savings, in *Factory* environment. For the Gaussian method, detecting the PC stage significantly outperforms the single-stage detection and recovery in all environments. For the autoencoder-based technique, detecting the PC stage achieves slightly better performance than only detecting control in *Factory*, *Farm*, and *Sparse* environment. However, in *Dense* environment, detecting the PC stage with the autoencoder-based scheme greatly outperforms detecting the control stage by 47.4%. Results indicate that a UAV achieves similar or higher performance by monitoring more stages, and the performance benefit is greater for complex environments.

D. Compute Overhead

Software-Level Protection: We study the overhead of the proposed software-level anomaly detection and recovery scheme across the tested environments. The overhead is the total detection and recomputation time for each mission. Table V shows that the overall overhead of the autoencoder is much smaller than the Gaussian-based technique. The overhead of the Gaussian-based technique is dominated by the recovery of perception and planning stages, which is

TABLE V
COMPUTE TIME OVERHEAD OF DETECTION AND RECOVERY

Environment	Factory		Farm		Sparse		Dense	
	DET	RECOV	DET	RECOV	DET	RECOV	DET	RECOV
Perception	<0.0001%	0.9603%	<0.0001%	1.0902%	<0.0001%	0.9788%	<0.0001%	1.1932%
Planning	<0.0001%	1.0199%	<0.0001%	0.7801%	<0.0001%	0.9421%	<0.0001%	1.0279%
Control	0.0008%	<0.0001%	0.0007%	<0.0001%	0.0009%	<0.0001%	0.0012%	<0.0001%
sum (Gaussian)	1.9810%		1.8710%		1.9218%		2.2223%	
PPC	0.0042%	<0.0001%	0.0037%	<0.0001%	0.0047%	<0.0001%	0.0062%	<0.0001%
sum (AutoE)	0.0042%		0.0037%		0.0047%		0.0062%	

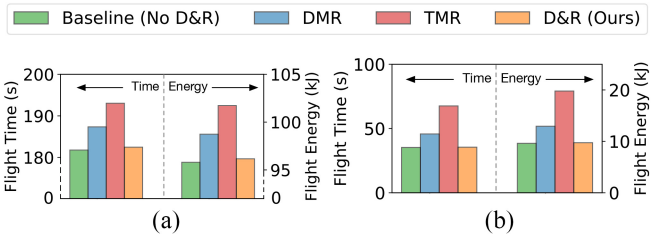

Fig. 15. Comparison of DMR, TMR, and the anomaly detection and recovery schemes on ARM Cortex-A57. (a) AirSim UAV. (b) DJI Spark.

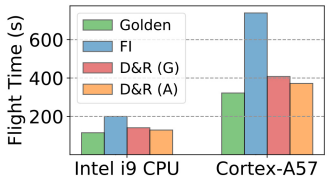

	i9-9940X	Cortex-A57
Core Number	14	4
Core Freq. (GHz)	3.3	2
Power (Watt)	165	<15
Flight time (s)	115	322
Flight energy (kJ)	61.7	177.1

Fig. 16. Comparison of detection and recovery schemes.

around 289 ms for each occupancy map generation and 83 ms for each trajectory generation. On the other hand, even if the autoencoder-based technique's detection overhead is higher, the recovery overhead is negligible as the control stage recomputation only takes 0.46 ms. As the scheme is operated at the software level with negligible overhead, it is possible to deploy multiple anomaly detection nodes to improve the robustness of detection nodes.

Hardware-Level Protection: To demonstrate the benefits of our schemes over redundancy-based hardware protections, we adopt a UAV visual performance model from [59] to evaluate the performance overhead of microarchitecture-based redundancy schemes (DMR and TMR) on UAV. Two types of UAVs, AirSim UAV and DJI Spark (with the same specs as [60]), are used as experimental platforms. Fig. 15 shows that TMR incurs a flight time increase by 1.06× on AirSim UAV and 1.91× on DJI compared to the anomaly detection scheme. The rationale is that hardware redundancy brings higher compute power with higher thermal design power and weight, thus lowering flight velocity and increasing flight time. Given the tight resource constraints of the UAV system, our scheme demonstrates negligible performance overhead.

E. Computing Platform Comparison

To show the portability we conduct FI on different platforms by introducing a single bit-flip at the inter-kernel states as in Section IV-C. Fig. 16 shows a similar error trend for both platforms. On the TX2, the worst flight time increases 2.8× since TX2 is an edge platform that has slower responses to environmental changes. However, with the anomaly detection ROS node continuously monitoring the anomaly of inter-kernel states, the flight time is recovered by 79.3% and 88.0% with Gaussian-based and autoencoder-based techniques.

VIII. CONCLUSION AND FUTURE WORK

Safety is paramount in autonomous vehicles. We present the first ROSFI fault analysis framework to enable system-level resilience analysis and show that system-level analysis is essential to capture system vulnerability compared to isolated, single-kernel FI analysis which is the common approach. Furthermore, we propose two anomaly detection and recovery schemes and demonstrate that with less than 0.0062% compute overhead, the autoencoder-based scheme can recover up to 100% failure cases in the tested scenario. Being an instruction-level FI framework, ROSFI has limitations. Nonetheless, we believe it moves the field of fault and resilience analysis forward in a significant way, providing a unique contribution of application-aware fault and resilience analysis in the context of robotics. Future directions include extending the fault model to consider microarchitecture-level errors that also manifest as bit flips in architecture states read by an instruction and incorporating different AV pipelines.

REFERENCES

[1] G. Li et al., "Understanding error propagation in deep learning neural network (DNN) accelerators and applications," in *Proc. Int. Conf. High Perform. Comput., Netw., Storage Anal.*, 2017, pp. 1–12.
[2] H. D. Dixit et al., "Silent data corruptions at scale," 2021, *arXiv:2102.11245*.
[3] P. H. Hochschild et al., "Cores that don't count," in *Proc. Workshop Hot Topics Oper. Syst.*, 2021, pp. 9–16.
[4] S. Jha et al., "ML-based fault injection for autonomous vehicles: A case for Bayesian fault injection," in *Proc. 49th Annu. IEEE/IFIP Int. Conf. Dependable Syst. Netw. (DSN)*, 2019, pp. 112–124.
[5] S. Jha, S. Cui, S. S. Banerjee, T. Tsai, Z. Kalbarczyk, and R. Iyer, "ML-driven malware that targets AV safety," 2020, *arXiv:2004.13004*.
[6] H. Shakhatreh et al., "Unmanned aerial vehicles (UAVs): A survey on civil applications and key research challenges," *IEEE Access*, vol. 7, pp. 48572–48634, 2019.
[7] Z. Wan et al., "A survey of FPGA-based robotic computing," *IEEE Circuits Syst. Mag.*, vol. 21, no. 2, pp. 48–74, 2nd Quart., 2021.
[8] S. S. Mukherjee, J. Emer, and S. K. Reinhardt, "The soft error problem: An architectural perspective," in *Proc. 11th Int. Symp. High-Perform. Comput. Archit.*, 2005, pp. 243–247.
[9] R. Bertran et al., "Voltage noise in multi-core processors: Empirical characterization and optimization opportunities," in *Proc. 47th Annu. IEEE/ACM Int. Symp. Microarchit.*, 2014, pp. 368–380.
[10] Z. Wan, K. Swaminathan, P.-Y. Chen, N. Chandramoorthy, and A. Raychowdhury, "Analyzing and improving resilience and robustness of autonomous systems," in *Proc. 41st IEEE/ACM Int. Conf. Comput.-Aided Design*, 2022, pp. 1–9.
[11] Y.-S. Hsiao et al., "MAVFI: An end-to-end fault analysis framework with anomaly detection and recovery for micro aerial vehicles," in *Proc. Des., Autom. Test Eur. Conf. Exhibit. (DATE)*, 2023, pp. 1–6.
[12] R. Nathan and D. J. Sorin, "Nostradamus: Low-cost hardware-only error detection for processor cores," in *Proc. Des., Autom. Test Eur. Conf. Exhibit. (DATE)*, 2014, pp. 1–6.
[13] E. Talpes et al., "Compute solution for Tesla's full self-driving computer," *IEEE Micro*, vol. 40, no. 2, pp. 25–35, Mar./Apr. 2020.
[14] S. K. S. Hari, M. Sullivan, T. Tsai, and S. W. Keckler, "Making convolutions resilient via algorithm-based error detection techniques," *IEEE Trans. Dependable Secure Comput.*, vol. 19, no. 4, pp. 2546–2558, Jul./Aug. 2022.
[15] S. K. S. Hari, T. Tsai, M. Stephenson, S. W. Keckler, and J. Emer, "SASSIFI: An architecture-level fault injection tool for GPU application resilience evaluation," in *Proc. IEEE Int. Symp. Perform. Anal. Syst. Softw. (ISPASS)*, 2017, pp. 249–258.

[16] G. Li, K. Pattabiraman, S. K. S. Hari, M. Sullivan, and T. Tsai, "Modeling soft-error propagation in programs," in *Proc. 48th Annu. IEEE/IFIP Int. Conf. Dependable Syst. Netw. (DSN)*, 2018, pp. 27–38.
[17] Z. Chen, G. Li, K. Pattabiraman, and N. DeBardeleben, "BinFI: An efficient fault injector for safety-critical machine learning systems," in *Proc. Int. Conf. High Perform. Comput., Netw., Storage Anal.*, 2019, pp. 1–23.
[18] G. Papadimitriou and D. Gizopoulos, "Demystifying the system vulnerability stack: Transient fault effects across the layers," in *Proc. ACM/IEEE 48th Annu. Int. Symp. Comput. Archit. (ISCA)*, 2021, pp. 902–915.
[19] Y.-S. Hsiao et al., "Zhuyi: Perception processing rate estimation for safety in autonomous vehicles," in *Proc. 59th ACM/IEEE Design Automat. Conf.*, 2022, pp. 289–294.
[20] Y. Gan, P. Whatmough, J. Leng, B. Yu, S. Liu, and Y. Zhu, "Braum: Analyzing and protecting autonomous machine software stack," in *Proc. IEEE 33rd Int. Symp. Software Rel. Eng. (ISSRE)*, 2022, pp. 85–96,
[21] R. Palin, D. Ward, I. Habli, and R. Rivett, "ISO 26262 safety cases: Compliance and assurance," in *Proc. 6th IET Int. Conf. Syst. Safety*, 2011, pp. 1–6.
[22] G. Li et al., "AV-FUZZER: Finding safety violations in autonomous driving systems," in *Proc. IEEE 31st Int. Symp. Softw. Rel. Eng. (ISSRE)*, 2020, pp. 25–36.
[23] Z. Wan, A. Anwar, Y.-S. Hsiao, T. Jia, V. J. Reddi, and A. Raychowdhury, "Analyzing and improving fault tolerance of learning-based navigation systems," in *Proc. 58th ACM/IEEE Design Autom. Conf. (DAC)*, 2021, pp. 841–846.
[24] Z. Wan et al., "FRL-FI: Transient fault analysis for federated reinforcement learning-based navigation systems," in *Proc. Des., Autom. Test Eur. Conf. Exhibit. (DATE)*, 2022, pp. 430–435.
[25] Z. Wan, N. Chandramoorthy, K. Swaminathan, P.-Y. Chen, V. J. Reddi, and A. Raychowdhury, "BERRY: Bit error robustness for energy-efficient reinforcement learning-based autonomous systems," in *Proc. 60th ACM/IEEE Design Autom. Conf. (DAC)*, 2023, pp. 1–6.
[26] M. Quigley et al., "ROS: An open-source robot operating system," in *Proc. ICRA Workshop Open Source Softw.*, vol. 3, 2009, p. 5.
[27] V. Mayoral-Vilches et al., "RobotPerf: An open-source, vendor-agnostic, benchmarking suite for evaluating robotics computing system performance," 2023, *arXiv:2309.09212*.
[28] B. Boroujerdian, H. Genc, S. Krishnan, W. Cui, A. Faust, and V. Reddi, "MAVBench: Micro aerial vehicle benchmarking," in *Proc. 51st Annu. IEEE/ACM Int. Symp. Microarchit. (MICRO)*, 2018, pp. 894–907.
[29] S. Liu, Z. Wan, B. Yu, and Y. Wang, "Robotic computing on FPGAs," in *Synthesis Lectures Computer Architecture*, vol. 16, Cham, Switzerland: Springer, 2021, pp. 1–218.
[30] Q. Liu, Z. Wan, B. Yu, W. Liu, S. Liu, and A. Raychowdhury, "An energy-efficient and runtime-reconfigurable fpga-based accelerator for robotic localization systems," in *Proc. IEEE Custom Integr. Circuits Conf. (CICC)*, 2022, pp. 01–02.
[31] R. B. Rusu and S. Cousins, "3D is here: Point cloud library (PCL)," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2011, pp. 1–4.
[32] F. Fleuret, J. Berclaz, R. Lengagne, and P. Fua, "Multicamera people tracking with a probabilistic occupancy map," *IEEE Trans. pattern Anal. Mach. Intell.*, vol. 30, no. 2, pp. 267–282, Feb. 2008.
[33] D. Gonzalez, J. Perez, V. Milanes, and F. Nashashibi, "A review of motion planning techniques for automated vehicles," *IEEE Trans. Intell. Transp. Syst.*, vol. 17, no. 4, pp. 1135–1145, Apr. 2016.
[34] K. H. Ang, G. Chong, and Y. Li, "PID control system analysis, design, and technology," *IEEE Trans. control Syst. Technol.*, vol. 13, no. 4, pp. 559–576, Jul. 2005.
[35] M. Maniatakos, N. Karimi, C. Tirumurti, A. Jas, and Y. Makris, "Instruction-level impact analysis of low-level faults in a modern microprocessor controller," *IEEE Trans. Comput.*, vol. 60, no. 9, pp. 1260–1273, Sep. 2011.
[36] N. Binkert et al., "The gem5 simulator," *ACM SIGARCH Comput. Archit. News*, vol. 39, no. 2, pp. 1–7, May 2011.
[37] A. Pellegrini et al., "Crashtest'ing swat: Accurate, gate-level evaluation of symptom-based resiliency solutions," in *Proc. Des., Autom. Test Eur. Conf. Exhibit. (DATE)*, 2012, pp. 1106–1109.
[38] M. Daněk, L. Kafka, L. Kohout, J. Sỳkora, and R. Bartosiński, "The leon3 processor," in *UTLEON3: Exploring Fine-Grain Multi-Threading in FPGAs*. New York, NY, USA: Springer, 2013, pp. 9–14.
[39] K. S. Yim, Z. Kalbarczyk, and R. K. Iyer, "Measurement-based analysis of fault and error sensitivities of dynamic memory," in *Proc. IEEE/IFIP Int. Conf. Dependable Syst. Netw. (DSN)*, 2010, pp. 431–436.
[40] S. E. Michalak et al., "Assessment of the impact of cosmic-ray-induced neutrons on hardware in the roadrunner supercomputer," *IEEE Trans. Device Mater. Rel.*, vol. 12, no. 2, pp. 445–454, Jun. 2012.
[41] H. Cho, S. Mirkhani, C.-Y. Cher, J. A. Abraham, and S. Mitra, "Quantitative evaluation of soft error injection techniques for robust system design," in *Proc. 50th Annu. Des. Autom. Conf.*, 2013, pp. 1–10.
[42] S. Shah, D. Dey, C. Lovett, and A. Kapoor, "AirSim: High-fidelity visual and physical simulation for autonomous vehicles," 2017, *arXiv:1705.05065*.
[43] J. Wei, A. Thomas, G. Li, and K. Pattabiraman, "Quantifying the accuracy of high-level fault injection techniques for hardware faults," in *Proc. 44th Annu. IEEE/IFIP Int. Conf. Dependable Syst. Netw.*, 2014, pp. 375–382.
[44] Q. Lu, M. Farahani, J. Wei, A. Thomas, and K. Pattabiraman, "LLFI: An intermediate code-level fault injection tool for hardware faults," in *Proc. IEEE Int. Conf. Softw. Qual., Rel. Security*, 2015, pp. 11–16.
[45] A. Mahmoud et al., "Minotaur: Adapting software testing techniques for hardware errors," in *Proc. 24th Int. Conf. Archit. Support Program. Lang. Oper. Syst.*, 2019, pp. 1087–1103.
[46] Q. Lu, G. Li, K. Pattabiraman, M. S. Gupta, and J. A. Rivers, "Configurable detection of SDC-causing errors in programs," *ACM Trans. Embed. Comput. Syst. (TECS)*, vol. 16, no. 3, pp. 1–25, Mar. 2017.
[47] B. Sangchoolie, K. Pattabiraman, and J. Karlsson, "One bit is (not) enough: An empirical study of the impact of single and multiple bit-flip errors," in *Proc. 47th Annu. IEEE/IFIP Int. Conf. Dependable Syst. Netw. (DSN)*, 2017, pp. 97–108.
[48] E. Cheng et al., "CLEAR: Cross-layer exploration for architecting resilience: Combining hardware and software techniques to tolerate soft errors in processor cores," in *Proc. 53rd Annu. Des. Autom. Conf.*, 2016, pp. 1–6.
[49] V. Mayoral-Vilches, S. M. Neuman, B. Plancher, and V. J. Reddi, "RobotCore: An open architecture for hardware acceleration in ROS 2," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2022, pp. 1–8.
[50] "NVBitFI: An architecture-level fault injection tool for GPU application resilience evaluations." github.com. [Online]. Available: https://github.com/NVlabs/nvbitfi
[51] V. Porpodas, "ZOFI: Zero-overhead fault injection tool for fast transient fault coverage analysis," 2019, *arXiv:1906.09390*.
[52] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "OctoMap: An efficient probabilistic 3D mapping framework based on octrees," *Auton. robots*, vol. 34, pp. 189–206, 2013.
[53] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *Int. J. Robot. Res.*, vol. 30, no. 7, pp. 846–894, 2011.
[54] B. Fang, Q. Lu, K. Pattabiraman, M. Ripeanu, and S. Gurumurthi, "ePVF: An enhanced program vulnerability factor methodology for cross-layer resilience analysis," in *Proc. 46th Annu. IEEE/IFIP Int. Conf. Dependable Syst. Netw. (DSN)*, 2016, pp. 168–179.
[55] F. F. Dos Santos et al., "Analyzing and increasing the reliability of convolutional neural networks on GPUs," *IEEE Trans. Rel.*, vol. 68, no. 2, pp. 663–677, Jun. 2019.
[56] L. Ruff et al., "A unifying review of deep and shallow anomaly detection," *Proc. IEEE*, vol. 109, no. 5, pp. 756–795, May 2021.
[57] *Open Sources Robotics Foundation*. ros. Accessed: Jan. 24, 2022. [Online]. Available: http://wiki.ros.org/roslaunch/XML/node
[58] D. E. Knuth, *Art of Computer Programming, Volume 2: Seminumerical Algorithms*. Boston, MA, USA: Addison-Wesley Prof., 2014.
[59] S. Krishnan et al., "The sky is not the limit: A visual performance model for cyber-physical co-design in autonomous machines," *IEEE Comput. Archit. Lett.*, vol. 19, no. 1, pp. 38–42, Jan.–Jun. 2020.
[60] S. Krishnan, Z. Wan, K. Bhardwaj, N. Jadhav, A. Faust, and V. J. Reddi, "Roofline model for UAVs: A bottleneck analysis tool for onboard compute characterization of autonomous unmanned aerial vehicles," in *Proc. IEEE Int. Symp. Perform. Anal. Syst. Softw. (ISPASS)*, 2022, pp. 162–174.

Yu-Shun Hsiao received the B.E. degree in electrical engineering from National Tsing Hua University, Hsinchu, Taiwan, in 2018. He is currently pursuing the Ph.D. degree in computer science with Harvard University, Cambridge, MA, USA.

He was a Research Scientist Intern with NVIDIA Architecture Research Team, Cambridge, in 2021 and 2022. His research interests include computer architecture and system-software design to enable efficient and resilient autonomous machines, including vehicles, drones, and robotic arms.

Zishen Wan (Student Member, IEEE) received the B.E. degree in electrical engineering and automation from the Harbin Institute of Technology, Harbin, China, in 2018, and the M.S. degree in electrical engineering from Harvard University, Cambridge, MA, USA, in 2020. He is currently pursuing the Ph.D. degree with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA.

His research interests include computer architecture, VLSI, and embedded systems, with a focus on designing efficient and reliable hardware and systems for autonomous machines and edge intelligence.

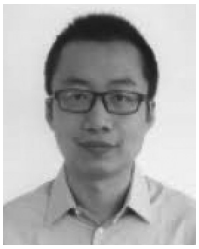

Tianyu Jia (Member, IEEE) received the B.S. and M.S. degrees from the Beijing University of Posts and Telecommunications, Beijing, China, in 2011 and 2014, respectively, and the Ph.D. degree in computer engineering from Northwestern University, Evanston, IL, USA, in 2019.

He was a Postdoctoral Fellow with Harvard University, Cambridge, MA, USA, and an Assistant Research Professor with Carnegie Mellon University, Pittsburgh, PA, USA. He is currently an Assistant Professor with Peking University, Beijing. His research interests include accelerator design for domain-specific applications, heterogeneous SoC design, and optimization.

Dr. Jia was awarded the IEEE Solid-State Circuit Society Predoctoral Achievement Award in 2020.

Radhika Ghosal received the B.Tech. degree in computer science and engineering from the Indraprastha Institute of Information Technology, New Delhi, India, in 2019. She is currently pursuing the Ph.D. degree in computer science with Harvard University, Cambridge, MA, USA.

Her research is supported by the NSF Graduate Research Fellowship. Her research interests lie in designing high-performance compute hardware and software for robotics.

Abdulrahman Mahmoud (Member, IEEE) received the B.S.E. degree in electrical engineering from Princeton University, Princeton, NJ, USA, in 2013, and the Ph.D. degree in computer science from the University of Illinois at Urbana–Champaign, Urbana, IL, USA, in 2020.

He is currently a Postdoctoral Fellow with Harvard University, Cambridge, MA, USA. He is the lead developer of multiple research tools along this line of work, including PyTorchFI and GoldenEye, which have been well received and used for efficient ML design in recent years. His current research interests include resilient and efficient machine learning systems using hardware–software co-design.

Arijit Raychowdhury (Fellow, IEEE) received the Ph.D. degree in electrical and computer engineering from Purdue University, West Lafayette, IN, USA, in 2007.

He is currently the Steve W Chaddick Chair and a Professor with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA. He is also the Director of the Center for the Co-Design of Cognitive Systems (CoCoSys), a Joint University Microelectronics Program 2.0. His research interests include low-power digital and mixed-signal circuit design, signal processors, and exploring interactions of circuits with device technologies.

Dr. Raychowdhury is the winner of several awards, including the SRC Technical Excellence Award in 2021, the Qualcomm Faculty Award in 2021 and 2020, the IEEE/ACM Innovator under 40 Award, and the NSF CISE Research Initiation Initiative Award (CRII) in 2015. He is currently a Distinguished Lecturer of the IEEE Solid-State Circuits Society. He serves on the technical program committee of key circuits and design conferences, including ISSCC, VLSI Symposium, DAC, and CICC.

David Brooks (Fellow, IEEE) received the B.S. degree in electrical engineering from the University of Southern California, Los Angeles, CA, USA, in 1997, and the M.A. and Ph.D. degrees in electrical engineering from Princeton University, Princeton, NJ, USA, in 1999 and 2001, respectively.

He is currently the Haley Family Professor of Computer Science with the School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA. His current research interests include resilient and power-efficient computer hardware and software design for high-performance and embedded systems.

Dr. Brooks was a recipient of several honors and awards, including the ACM Maurice Wilkes Award and the ISCA Influential Paper Award.

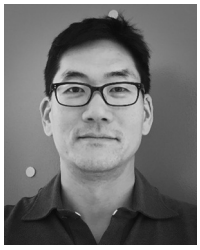

Gu-Yeon Wei (Senior Member, IEEE) received the B.S., M.S., and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, USA, in 1994, 1997, and 2001, respectively.

He is currently a Robert and Suzanne Case Professor of Electrical Engineering and Computer Science with the Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA. His research interests span multiple layers of a computing system: mixed-signal integrated circuits, computer architecture, and design tools for efficient hardware. His research efforts focus on identifying synergistic opportunities across these layers to develop energy-efficient solutions for a broad range of systems from flapping-wing microrobots to machine learning hardware for the Internet of Things devices to large-scale servers.

Vijay Janapa Reddi (Member, IEEE) received the Ph.D. degree in computer science from Harvard University, Cambridge, MA, USA, in 2010.

He is currently an Associate Professor with the John A. Paulson School of Engineering and Applied Sciences, Harvard University, where he is the Director of the Edge Computing Laboratory. His research interests include computer architecture and system software design, with an emphasis on the context of mobile and edge computing platforms based on machine learning.

Dr. Reddi was a recipient of multiple awards, including the Best Paper Award at MICRO 2005 and HPCA 2009, the IEEE's Top Picks in Computer Architecture Awards in 2006, 2010, 2011, 2016, and 2017, the Google Faculty Research Awards in 2012, 2013, 2015, and 2017, the Intel Early Career Award in 2013, the National Academy of Engineering Gilbreth Lecturer Honor in 2016, and the IEEE TCCA Young Computer Architect Award in 2016.