

Protecting Actuators in Safety-Critical IoT Systems from Control Spoofing Attacks

Monowar Hasan

Dept. of Computer Science, University of Illinois
mhasan11@illinois.edu

Sibin Mohan

Dept. of Computer Science, University of Illinois
sibin@illinois.edu

ABSTRACT

In this paper, we propose a framework called Contego-TEE to secure Internet-of-Things (IoT) edge devices with timing requirements from control spoofing attacks where an adversary sends malicious control signals to the actuators. We use a trusted computing base available in commodity processors (such as ARM TrustZone) and propose an invariant checking mechanism to ensure the security and safety of the physical system. A working prototype of Contego-TEE was developed using embedded Linux kernel. We demonstrate the feasibility of our approach for a robotic vehicle running on an ARM-based platform.

1 INTRODUCTION

Today's embedded and cyber-physical systems are ubiquitous. A large number of critical cyber-physical systems (*e.g.*, autonomous cars, drones, manufacturing systems, power grids, industrial control systems, *etc.*) have real-time (RT) properties (*e.g.*, strict timing and safety requirements). The current trend is to connect embedded RT devices to the Internet (*e.g.*, remote surveillance over wired/wireless network, connected vehicles through cellular wireless networks, *etc.*) and this gives rise to the real-time Internet-of-Things (RT-IoT) [1]. RT-IoT systems are intended to provide better user experience through stronger connectivity and better use of next-generation embedded devices, albeit with safety-critical properties. RT-IoT systems are also increasingly becoming targets for cyber-attacks. A number of high-profile attacks on RT-IoT systems, *e.g.*, denial-of-service (DoS) attacks mounted from IoT devices [2], Stuxnet [3], attack demonstrations by researchers on medical devices [4] and automobiles [5] have shown that the threat is real. Successful cyber attacks against such systems could lead to problems more serious than just loss of data or availability because of their critical nature [1]. Enabling security in RT-IoT, however, is often more challenging than generic IoT due to additional timing/safety constraints imposed by RT-enabled systems.

Since RT-IoT systems are largely based on sensing and actuation, any false/spoofed command to the actuators can disrupt the normal operation of the physical plant. Commonly used open-source RT-IoT development stacks (such as Linux) do not provide explicit control over actuation signals. For instance, if the application task obtains permission (say, root or other privileged user access) to the peripheral interface (*e.g.*, I2C [6]), it is possible to send arbitrary signals to the actuators. Let us consider an industrial robotic arm (running an embedded variant of Linux in an ARM Cortex-A53 platform [7]) that periodically opens and closes the grip to drop off and pick up objects in an assembly line. The movement of the grip is controlled by a servo. We use an open-source implementation [8] for this robotic arm where each operation is represented by a pulse value x (where $x = 577$ for `grip_open()` and $x = 420$ for `grip_close()`)

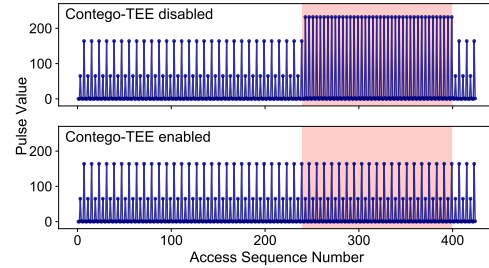


Figure 1: Demonstration of a control spoofing attack on a robotic control arm running embedded Linux.

and each pulse sends the following four 1 byte command sequences to the servo registers: $0 \& 0xFF$, $0 \gg 8$, $x \& 0xFF$, $x \gg 8$. An example of a spoofing attack for this control arm is presented in Fig. 1 (x -axis is the servo access sequence number and y -axis is the corresponding pulse value). Without any actuation command validation, it is possible to send arbitrary (high) pulses to the servo registers that prevents the grip from picking up/dropping objects (showing in the shaded region, see the top figure) that is not otherwise possible when our scheme (called Contego-TEE, see §3 for details) is enabled (bottom figure).

Our proposed framework, Contego-TEE, prevents the sending of malicious/undesired commands to physical actuators and ensures safety of the system. Specifically, we use the concept of Trusted Execution Environments (TEEs) [9] available in commodity processors (*e.g.*, ARM TrustZone [10], Intel SGX [11]) to ensure that our protection mechanisms can not be disabled even if the host OS is compromised. We develop a rule-based invariant checking and access control mechanism as well as design-time (schedulability) tests to ensure timing and safety requirements of the system. Contego-TEE specifically designed for *legacy systems* developed with Commodity-Off-The-Shelf (COTS) components and *does not require any modification to the application code/logic*.

In this paper we present the following contributions.

- A new framework called Contego-TEE to secure COTS-based RT-IoT systems against attacks that spoof control signals (§3.1).
- A runtime, rule-based invariant checking mechanism as well as design-time analysis to ensure security (and safety) of the physical plant (§3.2).
- An open source implementation and patch to the (embedded) Linux kernel that includes the Contego-TEE functionality (§4.1).

We use ARM TrustZone as a TEE and implement our solution in an ARM Cortex-A53 board (*i.e.*, Raspberry Pi [7]). We also

demonstrate the viability of our approach using a COTS rover platform (§4.2).

2 SYSTEM AND ADVERSARY MODEL

In the following, we first present background on RT-IoT systems and give an overview of a TEE-based architecture (e.g., ARM TrustZone). We then introduce our system model (§2.2) and describe our assumptions on the adversarial capabilities (§2.3).

2.1 Preliminaries

2.1.1 RT-IoT Systems: RT-IoT systems comprise IoT edge devices with RT capabilities. RT systems are those that, apart from a requirement for functional correctness, require that temporal properties be met as well. These temporal properties are often presented in the form of *deadlines*. The usefulness of results (say the performance of the actuators) produced by the system drops after the passage of a deadline. Some of the common properties and assumptions related to RT systems include [1]: (i) periodic/sporadic execution of set of tasks¹, (ii) strict timing and safety requirements, (iii) well-characterized execution time (e.g., execution times in the worst-case are known for all loops), (iv) limited resources (e.g., memory, processing power and energy).

RT-IoT systems are often designed based on the periodic task model [12], i.e., each task τ_i is characterized by a tuple: (C_i, T_i, D_i) where C_i is the Worst Case Execution Time (WCET), T_i is the period (e.g., inter-invocation time) and D_i is the deadline. Schedulability tests [13, 14] are used to determine if all tasks in the system meet their respective deadlines. If they do, then the taskset is deemed to be ‘*schedulable*’ and the system is considered *safe*.

2.1.2 TEE and ARM TrustZone: TEE is a set of hardware and software-based security extensions where the processors maintain a separated subsystem in addition to the traditional OS components. TEE technology has been implemented on commercial secure hardware such as ARM TrustZone [10] and Intel SGX [11]. In this work we consider TrustZone as the building block of Contego-TEE due to wide acceptability of ARM processors for embedded IoT systems – although our framework can be ported into other TEE platforms without loss of generality. ARM partnered with GlobalPlatform and has defined new TEE APIs [15]. TrustZone encompasses the following major features [16]: (a) safe and secure boot (to ensure all software components are in a trusted state before launching the OS); (b) isolated execution of critical applications (i.e., in a secure enclave) and (c) protection for trusted applications data (in terms of integrity and confidentiality).

ARM TrustZone contains two different privilege blocks: (i) *Normal World* (NW) and (ii) *Secure World* (SW). The NW is the untrusted environment running a commodity untrusted OS where SW is a protected computing block that only runs privileged instructions. SW in TrustZone defines the memory regions that can only be accessed by privileged instructions and the code that runs in the SW has higher privilege than the NW. Hardware logic ensures that the resources in the secure world can not be accessed from the normal world (e.g., if the code running in the NW tries to access protected memory regions, TrustZone throws a hardware

¹The ‘task’ in RT-IoT systems can trivially be mapped with the concept of *process* or *thread* in general-purpose OSes.

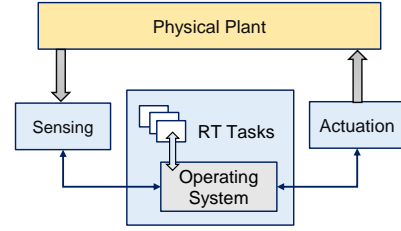


Figure 2: High-level schematic of a RT control system.

exception). The SW instructions are triggered when a specific flag in the processor e.g., Non-secure (NS) bit in the Secure Configuration Register (SCR) is not set. These two worlds bridge via a software module referred to as *Secure Monitor*. The context switch between the NW and SW is performed through a *Secure Monitor Call* (SMC).

In this work we use the TrustZone functionality to prevent the malicious commands from being sent to the actuators (See §3 for details). We now present our system and adversary model.

2.2 System Model

In Fig. 2 we present a high-level illustration of a RT control system. We consider a set of periodic RT control tasks $\Gamma = \{\tau_1, \tau_2, \dots, \tau_N\}$ that execute on single processor². The physical system consists of a set of M actuators (e.g., servo, motor, buzzer): $\{\pi_1, \pi_2, \dots, \pi_M\}$. RT tasks periodically issue commands to the actuators to control physical entities (e.g., wheel, propeller, alarm, robotic grip, etc.). We assume that each task is allowed to access a subset of peripherals. We represent this access permission as an $N \times M$ Boolean matrix $\mathbf{A} = [a_{ik}]$ where $a_{ik} = 1$ represents task τ_i can send commands to actuator π_k . We also assume that the RT tasks finish computation before their deadline, e.g., the tasks are schedulable³.

2.3 Adversary Model

We consider the following adversarial capabilities: (a) *Integrity Violation* – an adversary may insert a malicious task (that respects the RT guarantees) and/or modify exiting control logic to manipulate actuator commands and control system behavior in undesirable ways; (b) *Denial of Service (DoS)* – the attacker may take control of the RT task(s) and destabilize the physical plant e.g., by sending multiple control requests in a burst that may result in a malfunctioning actuator, or worse, damage the actual hardware/actuator and even threaten the safety of the system.

The attacker can gain privileged (e.g., root) access to perform adversarial actions (e.g., to spoof control signals). We do not make any assumptions as to how the compromised tasks enter the device. For instance, bad software engineering practices leave vulnerabilities in the systems [17]. When the system is developed using a multi-vendor model [18] (where its components are manufactured and integrated by different vendors) a malicious code logic may be injected (say by a less-trusted vendor) during deployment. The adversary may also induce end-users to download the modified source code, say by using social engineering

²Since majority of the RT-IoT edge devices still use single core chips due to simplicity and determinism.

³In the Appendix we present formal expressions to determine schedulability of the tasks.

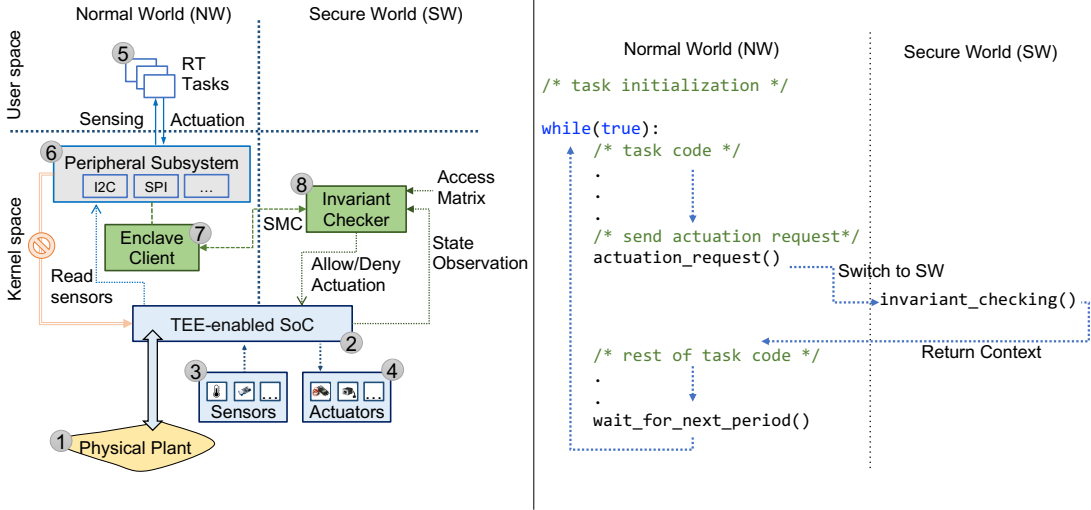


Figure 3: Overview of Contego-TEE system design (left) and high-level control flow of RT tasks in Contego-TEE (right).

tactics [19]. We also assume that the attackers do not have any physical access (e.g., they can not physically control/turn off/damage the actuators).

3 ACTUATION MONITORING FRAMEWORK

In the following we first introduce the Contego-TEE framework (§3.1). We then present mechanisms to detect any abnormal control commands issued by (rogue) tasks and analyze schedulability conditions that ensures our (invariant) checking techniques can be enforced at runtime (§3.2).

3.1 Overview and Architecture

As mentioned earlier, to secure RT-IoT platforms we propose a TEE-based architecture that monitors actuation commands sent to the physical entities. At the high-level, our design is based on the Simplex architecture [20]. Researchers use Simplex-based architecture for time-critical cyber-physical systems to provide fault-tolerance [21, 22] and recently, security [23–25]. A Simplex system consists of the following main components: (a) under normal operating conditions a *High-Performance (complex) Controller* actuates the physical plant (such a controller may be unverifiable due to its complexity, yet it must actuate a safety-critical system); (b) if, during operation the system state becomes unstable (e.g., it is in danger of violating a safety condition), a *Safety Controller* takes over and (c) the exact switching behavior is implemented by a *Decision Module* that decides which controller output will drive the plant. In our context, we use a trusted (and verified) computing module (this is analogous to the safety controller) executed in a secure enclave (viz., SW) and ensures that even if the (potentially untrusted) NW RT tasks are compromised, an adversary can not send false signals to the physical actuators.

In Fig. 3 we illustrate the high-level overview of Contego-TEE design and control flow of the RT tasks. Contego-TEE contains the following essential components: (a) a *TEE-enabled SoC* (System-on-Chip) such as those supported by ARM TrustZone [10] (block ② in the figure); (b) an *Enclave Client* (block ⑦) that is used to communicate between NW and SW and (c) an *Invariant Checker* (block ⑧) that is used to monitor (and validate) the actuation commands. The physical plant (①) is connected with sensors (②)

and actuators (③) and controlled by the (potentially vulnerable) RT tasks (⑤). RT tasks execute in untrusted NW and issue system calls (e.g., `read()`, `write()`, `ioctl()`) to access the sensors/actuators using specific interface such as I2C [6] and/or SPI [26]. Contego-TEE ensures that RT tasks cannot directly send any actuation commands (e.g., it breaks the bridge between ⑥, ② and ④). We do this by placing a dispatcher (e.g., enclave client) between the peripheral subsystem and actual hardware. As a result, before issuing any command to the physical actuators, it will be validated by our trusted application (e.g., invariant checker) running inside the secure enclave (i.e., in the SW). In particular, when a RT task τ_i sends an actuation command x_{ik}^t to any peripheral π_k at time t , enclave client traps those request and forwards the command to the invariant checker using SMC. Depending on the access permission matrix \mathbf{A} and current system state $\mathcal{S}(t)$, invariant checker then decides whether the given command x_{ik}^t can be issued to the actuator π_k (refer to §3.2 for details). In Contego-TEE, both the enclave client and invariant checker operate in the privileged mode (e.g., kernel space) so that it can directly control low-level hardware. By using the enclave client (to invoke context switching) and invariant checking mechanisms, Contego-TEE ensures that even if the NW RT tasks are compromised, an adversary can not send false signals to the actuators. We note that unlike NW RT tasks that may perform other computation, the invariant checker contains a small, verified, code blocks that is used to monitor only actuation requests. We also note that Contego-TEE *does not require any application-level modifications*, e.g., developers can execute unmodified, existing legacy RT tasks, using our Contego-TEE enabled OS-kernel (refer to §4.1 for implementation/porting details).

3.2 Invariant Checking and Timing Analysis

3.2.1 Invariant Checking: In order to validate each actuation command invoked by the RT tasks, Contego-TEE performs various actions. One obvious access control mechanism is to ensure that a task τ_i can access a given actuator π_k only if the task has the required permission (e.g., $a_{ik} = 1$). Contego-TEE therefore denies all the actuation commands from tasks if the corresponding access flag is zero. However, if the attacker can compromise a task with legitimate access (to a given set of actuators) then the (victim)

Table 1: Applicability of Contego-TEE for Various RT-IoT Platforms

Platform	Application Domain	Actuators	Possible Invariant Conditions*	Response	Remarks
Water/air monitoring system	Home/industrial automation	Buzzer, display	(a) Send high pulse to buzzer only if water-level is high/air quality abnormal/detect smoke; (b) do not display alert if the system state is normal	IGNORE	Ignore all commands that fail invariant checking
Surveillance system	Home/industrial automation	Servo, buzzer	(a) Trigger alarm only if there is an impact/object detected in camera; (b) rotate camera (using servos) only within allowable pan/tilt angle	IGNORE	Ignore all commands that fail invariant checking
Infusion/syringe pump	Health-care	Motor, display	(a) Drive the motor only to allowable positions/rates (b) display only the amount of fluid infused (e.g., obtained from motor encoders)	IGNORE	Ignore actuation when the task tries to infuse wrong amount of fluid
Robotic arm	Manufacturing	Servo, buzzer	(a) Check the servo pulse sequences matches with the desired (design-time) sequence; (b) do not raise alarm if the pulse sequence is normal	IGNORE, FAIL-SAFE	If mismatch, use the predefined sequence; ignore other pulses using rate-control rule
Robotic vehicle (aerial/ground)	Manufacturing, surveillance, agriculture	Servo, motor	(a) Check if the robot is following the mission; (b) allow only predefined number of actuation commands per period	IGNORE, FAIL-SAFE	Ignore command using rate-control rule. If it deviates from the mission, use predefined command and/or state-observations

*We omit mathematical expressions for readability.

task may send arbitrary commands to the actuators. Therefore in addition to checking access matrix \mathbf{A} , Contego-TEE also performs checking of system invariants and monitors the number of actuation commands for a given time interval as we discuss below.

State Invariant Checking: Invariant checking [27] is useful to detect control spoofing attacks. For a given RT-IoT platform we do this by considering the availability of an invariant checking function $\text{CheckInv}(\tau_i, \pi_k)$ that predicts the actuation signal and only allows access if the output of the function matches that of the requested command. In particular, if a task τ_i sends actuation command x_{ik}^t at time t to any peripheral π_k and the task has the required permission (i.e., $a_{ik} = 1$), $\text{CheckInv}(\tau_i, \pi_k)$ first obtains system state $\mathcal{S}(t)$ by observing a set of signals $S_i = \{s_1, s_2, \dots, s_{L_i}\}$ and decides whether x_{ik}^t is valid for current state $\mathcal{S}(t)$. For example, consider a warehouse water monitoring system where an alarm is triggered only if the water level of the tank (measured by the sensor s_{WL}) is higher than a predefined threshold (θ_{WL}) and/or the water temperature (s_{WT}) is not in expected range (i.e., $[\theta_{WT}^1, \theta_{WT}^2]$). We represent this as the following invariant rule: $\text{INV}_W :: (s_{WL} > \theta_{WL}) \vee (s_{WT} \notin [\theta_{WT}^1, \theta_{WT}^2]) \rightarrow x = \text{ON} : x = \text{OFF}$, e.g., Contego-TEE will only allow the sending of the high pulse (i.e., $x = \text{ON}$) to the alarm system (say a buzzer) only if the invariant conditions are satisfied. We note that since Contego-TEE operates at the kernel-level, it can directly access raw signals without any interaction of NW RT tasks or other (user space) libraries.

Rate Control: Note that since RT systems are deterministic by design, the (worst-case) number of actuation requests can be bounded at design time [28]. Therefore, if a task τ_i tries to access actuator(s) more than expected within a given time interval (e.g., T_i), it may be indication of a possible attack. In such cases Contego-TEE will limit subsequent access requests from τ_i and prevent the sending of actuation commands to the hardware. We enforce rate control using the following invariant rule: $\text{INV}_{ik}^{RC} :: \Delta_{ik}(w) < \widehat{\theta}_{ik} \rightarrow \text{CheckInv}(\tau_i, \pi_k) : \text{ignore}$, i.e., Contego-TEE ignores further actuation commands if the number of requests $\Delta_{ik}(\cdot)$ from any job of τ_i within the (relative) time window $w \in [0, T_i]$ is exceeded a design-time threshold $\widehat{\theta}_{ik}$. Such a rate control mechanism is

specially useful to defend against DoS attacks where an attacker sends multiple actuation commands in a burst (say to quickly change the speed of wheels/propellers in robotic ground/aerial vehicles, abruptly move robotic arms, falsely toggles buzzers, etc.) to disrupt normal operations of the system.

3.2.2 Response Mechanisms: When there exists a mismatch between output of the $\text{CheckInv}(\cdot)$ and the requested actuation commands, Contego-TEE makes use of the following strategies to keep the physical system safe.

- **IGNORE:** this strategy prevents the execution of any actuation commands requested by RT tasks. Hence, actuators will not receive any signals from Contego-TEE and will continue to operate using the last known (uncompromised) commands. Contego-TEE will also ignore commands if the task makes multiple requests in a short time window (e.g., by using rate control rule).

- **FAIL-SAFE:** while the IGNORE strategy ensures that actuators will not get any abnormal signals, ignoring actuation commands (for a long time) may not be acceptable for highly dynamic systems such as unmanned ground/aerial vehicles (e.g., it may crash). Therefore, Contego-TEE also allows operation of a FAIL-SAFE mode, i.e., if it finds any mismatch, it ignores the requests from RT tasks and sends the predetermined (and/or based on the output of $\text{CheckInv}(\cdot)$) commands to make the system safe/operational. As an example, if there is a sudden change in the propeller speed of a UAV, the FAIL-SAFE strategy sets a safe, predefined speed, based on the current state of the UAV.

Depending on the target system, both of the above strategies may be required to keep the physical system operational. We note that invariant checking and response mechanisms are application dependent. Contego-TEE provides flexibility for the system engineers to develop appropriate mechanisms depending on the application requirements. In Table 1 we summarize possible invariant conditions and response mechanisms that are applicable for various RT-IoT platforms – however, this is by no stretch meant to an exhaustive list.

3.2.3 Schedulability Analysis: In order to perform invariant checking and execute the response mechanisms at runtime, we

Table 2: Summary of the Implementation Platform

Artifact	Configuration
Platform	Broadcom BCM2837 (Raspberry Pi 3)
CPU	1.2 GHz 64-bit ARM Cortex-A53
Memory	1 Gigabyte
Operating System	Linux (NW), OP-TEE (SW)
Kernel version	Linux kernel 4.16.56, OP-TEE core 3.4
Peripheral interface	I2C
Boot parameters	dtparam=i2c_arm=on, dtparam=spi=on

need to ensure that our framework should not cause delays and the timing requirements of RT tasks are satisfied (e.g., they complete execution before deadline). We therefore develop design-time schedulability tests that ensure the taskset is schedulable (refer to the Appendix for details). For instance, the RT task τ_i is schedulable in Contego-TEE if the Worst Case Response Time (WCRT) R_i^{TEE} is less than deadline, i.e., $R_i^{TEE} = C_i^{TEE} + I_i^{TEE} \leq D_i$, where C_i^{TEE} is the task WCET (including the time for world switching and invariant checking) and I_i^{TEE} is the interference⁴ from other tasks. The taskset Γ is referred to as schedulable if all the tasks are schedulable, viz., $R_i^{TEE} \leq D_i, \forall \tau_i \in \Gamma$.

4 EVALUATION

In this section we first present the implementation details of Contego-TEE (§4.1) and then show the viability of our approach using a case-study on a robotic vehicle (§4.2). Table 2 summarizes the system configurations and implementation details.

4.1 System Implementation

We implemented a proof-of-concept prototype of Contego-TEE on Raspberry Pi 3 (RPi3) Model B [7] (equipped with 1.2 GHz 64-bit ARMv8 CPU and 1 GB RAM). We selected RPi3 as our implementation platform since (a) it supports ARM TrustZone and (b) previous research has shown feasibility of deploying multiple IoT-specific applications on RPi3 [16, 19, 29–31]. We developed Contego-TEE using the Open-Portable Trusted Execution Environment (OP-TEE) [32] software stack that uses GlobalPlatform TEE APIs [15] to provide TrustZone functionality. OP-TEE provides a minimal secure kernel (called OP-TEE core) that can be run in parallel with the NW OS (e.g., Linux). In particular, we used Ubuntu 18.04 filesystem with a 64-bit Linux kernel (version 4.16.56) as the NW OS and our invariant checker is running on OP-TEE secure kernel (version 3.4). The enclave client was statically built with the Linux kernel. In order to implement the enclave client, we extended the Linux TEE interface (`/linux/drivers/tee/`) and enabled SMC from Linux kernel space⁵. We implemented the invariant checker as an OP-TEE kernel-level trusted application⁶ (e.g., in `/optee_os/core/arch/arm/pta/`). In our current implementation Contego-TEE supports actuators that are controlled via the

⁴In RT scheduling theory, the term ‘interference’ refers to the amount of time (from release to deadline) the task τ_i is ready but can not be scheduled due to execution of other tasks.

⁵Since GlobalPlatform APIs only support SMC from user space.

⁶This is known as PTA (Pseudo Trusted Application) in OP-TEE terminology.

I2C interface. Specifically, we modified the built-in structure `i2cdev_fops` (e.g., in `/linux/drivers/i2c/i2c-dev.c`) with our enclave client functions that is then switch the control to the invariant checker (e.g., by using SMC). Our implementation code is available in a public repository [33].

4.2 Case-Study: Robotic Vehicle

We implemented Contego-TEE in a COTS rover (named GoPiGo2, manufactured by Dexter Industries [34]) that can be used in multiple IoT-specific applications such as remote surveillance, agriculture, manufacturing, etc. [35]. The rover is equipped with two optical encoders that are connected to the motors (e.g., actuator in this setup): it can turn left by switching off the right encoder and vice-versa. The detailed specifications of the rover are available on the vendor website [34].

4.2.1 Results. We first demonstrate how Contego-TEE can be used to protect such systems from actuation attacks and then measure the performance overheads.

Security Analysis: For the following experiments, we conducted a line following mission where the robot steered from an initial location to a target location by following a line. The controller task was running as a NW Linux application and executed vendor-provided PID (Proportional–Integral–Derivative) closed-loop control [36] to track the planned path using the data received from sensors. The rover used the following commands: `fwd()`, `lft()`, `rht()`, `st_sp(δ)` for navigating the rover forward/left/right and set the speed to δ , respectively, where each command sent a 5-byte value to the actuator registers (e.g., wheel encoders/motors) using the I2C interface. For this mission we defined the following three invariant conditions⁷ that were used to monitor control signals (e.g., `cmd`): $INV_1 :: s_{LF} < -\theta \rightarrow cmd = st_sp(\delta_1) \wedge rht()$, $INV_2 :: s_{LF} > \theta \rightarrow cmd = st_sp(\delta_1) \wedge lft()$ and $INV_3 :: s_{LF} \in [-\theta, \theta] \rightarrow cmd = st_sp(\delta_2) \wedge fwd()$ where s_{LF} was the readings from the sensor, $\theta = 2500$ was a vendor-provided threshold (e.g., to follow the line) and $\delta_1, \delta_2 \in [0, 255]$ were used to set the speed of the rover. We show the case where the access flag is set (e.g., $a_{ik} = 1$) since Contego-TEE will trivially deny requests if the corresponding flag is zero. In our experiments we used both the FAIL-SAFE and IGNORE (to enforce rate control) strategies. For each actuation signal, our invariant checker matches with the desired signal and choose the appropriate strategy as we present in the following.

In Fig. 4a we illustrate our invariant checking mechanism with FAIL-SAFE strategy. The x-axis of the figure shows the time (e.g., count of the controller job) and the y-axis is the total distance travelled by the rover (e.g., readings from the encoders). In order to demonstrate malicious behavior, we followed a strategy similar to that considered in prior work [24, 25, 35, 37]. In particular, during program execution, we injected a logic bomb (during the shaded region in Fig. 4a) and sent erroneous commands to the controller. In this case, during the control spoofing attack, the rover deviated from the mission (e.g., PID control loop) and falsely sent commands to turn off one of the motors. As a result, when Contego-TEE was not active, the rover was not following the line and the encoder

⁷We manually inspected the vendor-provided control code and translated them into invariant conditions.

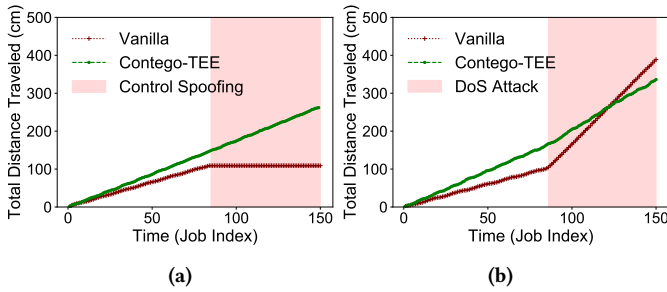


Figure 4: Illustration of Contego-TEE under (a) control spoofing and (b) DoS attacks. Contego-TEE prevents the sending of malicious commands to the motors and ensures that the rover moves at a steady speed.

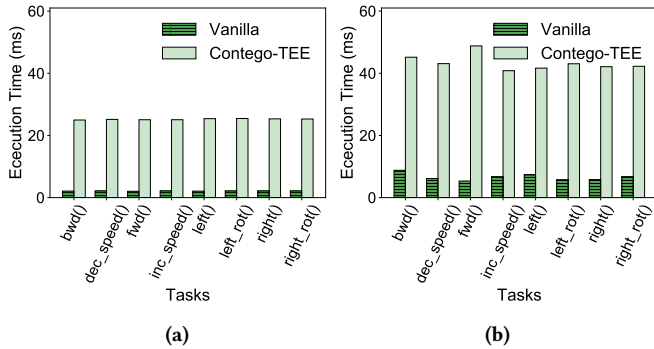


Figure 5: Runtime of rover control tasks with and without Contego-TEE: (a) for 99-th percentile and (b) worst-case. Contego-TEE increases the execution time by upto 43.47 ms (worst-case) and 23.31 ms (99th-percentile).

readings (*i.e.*, traversed distance) remained same (see the maroon line in the figure). We next executed the same code with Contego-TEE enabled (green curve in the figure). In this case, when each control command was issued, our checker followed the invariant conditions (*e.g.*, $INV_i, 1 \leq i \leq 3$) and sent desired commands to the motors (and hence the rover was moving as expected).

We next show the effect of our rate control mechanism (Fig. 4b). In this experiment, when the DoS logic bomb was triggered (shaded region in the figure) it sent multiple requests to increase the speed of the rover. When Contego-TEE was not enabled, this caused the rover to move faster and hence there was a rapid increase in the encoder readings (*e.g.*, maroon line, shaded region in the figure). In contrast, when Contego-TEE was active (green line), it disallowed multiple increase speed requests per period (*e.g.*, according to IGNORE strategy) and hence the rover followed the line with a steady speed.

Overhead Analysis: To measure the runtime overheads we conducted experiments with the vendor-provided control tasks [34] as a benchmark (Fig. 5). In this setup our invariant checker was following a rate control policy and ignored more than one actuation request per period (200 ms). The x-axis of Fig. 5 shows the control tasks and y-axis represents execution time (a) when Contego-TEE is not enabled (dark bar) and (b) with Contego-TEE enabled (light

bar). We present the timing results for 99th percentile (Fig. 5a) and worst-case (Fig. 5b). The timing values were measured using the Linux `clock_gettime()` system call with `CLOCK_MONOTONIC` clock parameter and we present data from 10,000 trials. As we see from the figure, Contego-TEE increases the execution time – this is expected due to (world) context switching as well for invariant checking. From our experiments we found that Contego-TEE increases execution times by (i) 34.11 to 43.47 ms (worst-case), (ii) 22.87 to 23.31 ms (99-th percentile) and (iii) 19.55 to 19.60 ms (average-case) for the various control tasks and hence can be used with 15 Hz (or slower) controllers (for this setup). This extra overhead results in increased security and we expect this could be acceptable for various RT-IoT platforms.

5 RELATED WORK

Enhancing security in time-critical cyber-physical systems is an active research area (see the related survey [1]). Perhaps the closest line of work to ours is PROTC [31] where a monitor in the SW enforces secure access control policy (given by the control center) for some peripherals of the drone and ensures that only authorized applications can access certain peripherals. Unlike our scheme, PROTC is limited for specific applications (*e.g.*, aerial robotic vehicles) and requires a centralized control center to validate/enforce security policies. In early work we proposed mechanisms to secure legacy time-critical systems [38–40]. Researchers also proposed anomaly detection approaches for robotic vehicles [35, 37, 41]. However these (prior) approaches do not provide any response mechanism and are vulnerable if the adversary can compromise the host OS.

There exist various hardware/software-based mechanisms and architectural frameworks [19, 23–25, 42, 43] to secure RT-IoT systems. However those frameworks are not designed to protect against control-specific attacks and may not be suitable for systems developed with COTS components. There also exist large number of research for generic IoT systems as well as use of TrustZone to secure traditional embedded/mobile applications (too many to enumerate here, refer to the related surveys [10, 44–46]) – however the consideration of time-critical and control-centric aspects of RT-IoT applications distinguish Contego-TEE from other research.

6 CONCLUSION

In this paper we presented a new framework named Contego-TEE that enhances the security and safety of the RT-IoT systems. We use a combination of trusted hardware, intrinsic real-time nature and domain-specific characteristics of such systems to detect control intrusions and prevent the physical plants from being misbehaved under attacks. We believe our framework is tangential and can be incorporated into multiple RT-IoT and cyber-physical domains.

REFERENCES

- [1] C.-Y. Chen, M. Hasan, and S. Mohan, “Securing real-time Internet-of-things,” *Sensors*, vol. 18, no. 12, 2018.
- [2] J. Westling, “Future of the Internet of things in mission critical applications,” 2016.
- [3] N. Falliere, L. O. Murchu, and E. Chien, “W32. stuxnet dossier,” *White paper, Symantec Corp., Security Response*, vol. 5, p. 6, 2011.
- [4] S. S. Clark and K. Fu, “Recent results in computer security for medical devices,” in *MobiHealth*, 2011, pp. 111–118.

- [5] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, T. Kohno *et al.*, "Comprehensive experimental analysis of automotive attack surfaces," in *USENIX Sec. Symp.*, 2011.
- [6] "i2c manual," Philips Semiconductors, 2003. [Online]. Available: <https://tinyurl.com/i2c-manual>
- [7] "Raspberry Pi," <https://www.raspberrypi.org/products/raspberry-pi-3-model-b/>.
- [8] "Robot arm control," <https://github.com/tutRPi/6DOF-Robot-Arm>.
- [9] M. Sabt, M. Achemlal, and A. Bouabdallah, "Trusted execution environment: What it is, and what it is not," in *IEEE Trustcom/BigDataSE/ISPA*, 2015, pp. 57–64.
- [10] S. Pinto and N. Santos, "Demystifying ARM TrustZone: A comprehensive survey," *ACM CSUR*, vol. 51, no. 6, p. 130, 2019.
- [11] V. Costan and S. Devadas, "Intel SGX Explained," *IACR Crypt. ePrint Arch.*, no. 086, pp. 1–118, 2016.
- [12] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard-real-time environment," *JACM*, vol. 20, no. 1, pp. 46–61, 1973.
- [13] N. Audsley, A. Burns, M. Richardson, K. Tindell, and A. J. Wellings, "Applying new scheduling theory to static priority pre-emptive scheduling," *SE Journal*, vol. 8, no. 5, pp. 284–292, 1993.
- [14] E. Bini and G. C. Buttazzo, "Schedulability analysis of periodic fixed priority systems," *IEEE Trans. on Comp.*, vol. 53, no. 11, pp. 1462–1473, 2004.
- [15] "TEE client API specification v1.0," <https://globalplatform.org/specs-library/tee-client-api-specification/>.
- [16] T. Liu, A. Hojjati, A. Bates, and K. Nahrstedt, "Alidrone: Enabling trustworthy proof-of-alibi for commercial drone compliance," in *IEEE ICDCS*, 2018, pp. 841–852.
- [17] F. Loi, A. Sivanathan, H. H. Gharakheili, A. Radford, and V. Sivaraman, "Systematically evaluating security and privacy for consumer IoT devices," in *ACM IoT&P*, 2017, pp. 1–6.
- [18] R. Pellizzoni, N. Paryab, M.-K. Yoon, S. Bak, S. Mohan, and R. B. Bobba, "A generalized model for preventing information leakage in hard real-time systems," in *IEEE RTAS*, 2015, pp. 271–282.
- [19] M.-K. Yoon, S. Mohan, J. Choi, M. Christodorescu, and L. Sha, "Learning execution contexts from system call distribution for anomaly detection in smart embedded system," in *ACM/IEEE IoTDI*, 2017, pp. 191–196.
- [20] L. Sha, "Using simplicity to control complexity," *IEEE Software*, vol. 18, no. 4, pp. 20–28, 2001.
- [21] X. Liu, Q. Wang, S. Gopalakrishnan, W. He, L. Sha, H. Ding, and K. Lee, "ORTEGA: An efficient and flexible online fault tolerance architecture for real-time control systems," *IEEE T. on Ind. Inf.*, vol. 4, no. 4, pp. 213–224, 2008.
- [22] X. Wang, N. Hovakimyan, and L. Sha, "L1Simplex: Fault-tolerant control of cyber-physical systems," in *2013 ACM/IEEE ICCPS*, 2013, pp. 41–50.
- [23] S. Mohan, S. Bak, E. Betti, H. Yun, L. Sha, and M. Caccamo, "S3A: Secure system simplex architecture for enhanced security and robustness of cyber-physical systems," in *ACM international conference on High confidence networked systems*. ACM, 2013, pp. 65–74.
- [24] F. Abdi, M. Hasan, S. Mohan, D. Agarwal, and M. Caccamo, "ReSecure: A restart-based security protocol for tightly actuated hard real-time systems," in *IEEE CERTS*, 2016, pp. 47–54.
- [25] M.-K. Yoon, S. Mohan, J. Choi, J.-E. Kim, and L. Sha, "SecureCore: A multicore-based intrusion detection architecture for real-time embedded systems," in *IEEE RTAS*, 2013, pp. 21–32.
- [26] "SPI block guide V04.01," Motorola Inc, 2004. [Online]. Available: <https://tinyurl.com/spi-block>
- [27] S. Adepu and A. Mathur, "From design to invariants: Detecting attacks on cyber physical systems," in *IEEE QRS-C*, 2017, pp. 533–540.
- [28] R. Wilhelm, J. Engblom, A. Ermedahl, N. Holsti, S. Thesing, D. Whalley, G. Bernat, C. Ferdinand, R. Heckmann, T. Mitra *et al.*, "The worst-case execution-time problem—overview of methods and survey of tools," *ACM TECS*, vol. 7, no. 3, p. 36, 2008.
- [29] L. Cheng, K. Tian, and D. D. Yao, "Orpheus: Enforcing cyber-physical execution semantics to defend against data-oriented attacks," in *ACM ACSAC*, 2017, pp. 315–326.
- [30] R. Liu and M. Srivastava, "VirtSense: Virtualize Sensing through ARM TrustZone on Internet-of-Things," in *ACM SysTEX*, 2018, pp. 2–7.
- [31] R. Liu and M. Srivastava, "PROTC: PROTeCting drone's peripherals through ARM trustzone," in *ACM Dronet*, 2017, pp. 1–6.
- [32] "Open Portable Trusted Execution Environment," <https://www.op-tee.org/>.
- [33] "Implementation code for Contego-TEE," https://github.com/mnwrhsn/rt_actuator_security.
- [34] "GoPiGo," <https://github.com/DexterInd/GoPiGo>.
- [35] P. Guo, H. Kim, N. Virani, J. Xu, M. Zhu, and P. Liu, "RoboADS: Anomaly detection against sensor and actuator misbehaviors in mobile robots," in *IEEE/IFIP DSN*, 2018, pp. 574–585.
- [36] "Dexter Industries Sensors," https://github.com/DexterInd/DI_Sensors.
- [37] H. Choi, W.-C. Lee, Y. Aafer, F. Fei, Z. Tu, X. Zhang, D. Xu, and X. Xinyan, "Detecting attacks against robotic vehicles: A control invariant approach," in *ACM CCS*, 2018, pp. 801–816.
- [38] M. Hasan, S. Mohan, R. Pellizzoni, and R. B. Bobba, "Contego: An adaptive framework for integrating security tasks in real-time systems," in *Euromicro ECRTS*, 2017, pp. 23:1–23:22.
- [39] M. Hasan, S. Mohan, R. B. Bobba, and R. Pellizzoni, "Exploring opportunistic execution for integrating security into legacy hard real-time systems," in *IEEE RTSS*, 2016, pp. 123–134.
- [40] M. Hasan, S. Mohan, R. Pellizzoni, and R. B. Bobba, "A design-space exploration for allocating security tasks in multicore real-time systems," in *DATE*, 2018, pp. 225–230.
- [41] F. Fei, Z. Tu, R. Yu, T. Kim, X. Zhang, D. Xu, and X. Deng, "Cross-layer retrofitting of UAVs against cyber-physical attacks," in *IEEE ICRA*, 2018, pp. 550–557.
- [42] M.-K. Yoon, S. Mohan, J. Choi, and L. Sha, "Memory heat map: anomaly detection in real-time embedded systems using memory behavior," in *ACM/EDAC/IEEE DAC*, 2015, pp. 1–6.
- [43] F. Abdi, C.-Y. Chen, M. Hasan, S. Liu, S. Mohan, and M. Caccamo, "Guaranteed physical security with restart-based design for cyber-physical systems," in *ACM/IEEE ICCPS*, 2018, pp. 10–21.
- [44] Y. Yang, L. Wu, G. Yin, L. Li, and H. Zhao, "A survey on security and privacy issues in Internet-of-Things," *IEEE IoT J.*, vol. 4, no. 5, pp. 1250–1258, 2017.
- [45] M. Ammar, G. Russello, and B. Crispo, "Internet of Things: A survey on the security of IoT frameworks," *Elsevier J. of Inf. Sec. & App.*, vol. 38, pp. 8–27, 2018.
- [46] W. Li, H. Chen, and H. Chen, "Research on ARM TrustZone," *ACM GetMobile*, vol. 22, no. 3, pp. 17–22, 2019.
- [47] M. Joseph and P. Pandya, "Finding response times in a real-time system," *The Comp. J.*, vol. 29, no. 5, pp. 390–395, 1986.

APPENDIX

Response Time Analysis for RT Tasks

Our schedulability test is based on the fixed-priority response time analysis proposed in RT literature [13]. Let N_i be the number of actuation request for τ_i and C_i^o is the additional computation time due to world switch and invariant checking. Then the WCET of τ_i can be represent as $C_i^{TEE} = C_i + N_i C_i^o$. Since our enclave client and invariant checker can serve one actuation request at a time (e.g., an atomic process), τ_i may be delayed due to processing requests of lower priority tasks. Let $B_i^{TEE} = \max_{\tau_j \in lp(\tau_i)} N_j C_j^o$ denote

the 'blocking' factor from tasks that are with lower-priority than τ_i (denoted as $lp(\tau_i)$). We note that the maximum computational demand for a given task τ_j in any interval length $0 \leq w \leq T_j$ can be no more than the maximum execution time required by one job of τ_j multiplied by the maximum number of jobs of τ_j that can execute in that interval [13, 14]. The maximum interference experience by τ_i from other tasks for an interval w can be expressed as: $B_i^{TEE} + \sum_{\tau_h \in hp(\tau_i)} \left\lceil \frac{w}{T_h} \right\rceil C_h^{TEE}$ where $hp(\tau_i)$ denotes the set of tasks with a priority higher than τ_i . Therefore, we can calculate the response time of τ_i (denoted as r_i) as follows:

$$r_i = C_i^{TEE} + B_i^{TEE} + \sum_{\tau_h \in hp(\tau_i)} \left\lceil \frac{r_i}{T_h} \right\rceil C_h^{TEE}. \quad (1)$$

The WCRT then can be obtained by solving this recurrence using an iterative fixed-point search, e.g., $R_i^{TEE} = r_i^{(\alpha)} = r_i^{(\alpha-1)}$ for some iteration α with initial condition $r_i^{(0)} = 0$. The iteration is guaranteed to be converged if we assume that the total processor utilization (i.e., $\sum_{\tau_i} \frac{C_i^{TEE}}{T_i}$) is less than 1 [47]. The taskset is considered

as 'unschedulable' if there exists an α such that $r_i^{(\alpha)} > D_i$. Such unschedulability result will hint the designers to update parameters (e.g., periods, number of tasks, invariant checking policies) to incorporate Contego-TEE framework for the target system.