

RESCUE: A Reconfigurable Scheduling Framework for Securing Multi-Core Real-Time Systems

ZAIN A. H. HAMMADEH, German Aerospace Center (DLR), Germany

MONOWAR HASAN, Washington State University, USA

MOHAMMAD HAMAD, Technical University of Munich, Germany

Modern real-time systems face increasing vulnerabilities to cyber-attacks, particularly those that use multi-core chips, where safety-critical and non-safety-critical tasks execute concurrently. Existing solutions for multicore systems often lack either determinism or cost-efficiency. This paper introduces an offline analysis technique that computes all feasible schedules for real-time tasks running on multi-core platforms. Our proposed technique isolates compromised tasks while ensuring a fail-operational system and supports low-cost, reconfigurable scheduling. The analytical models presented in this paper guarantee the hard real-time constraints of safety-critical tasks while allowing bounded deadline misses for some non-safety-critical tasks during an attack to enhance security. We name our scheme RESCUE. We conduct a comprehensive design-space exploration and evaluate its real-world efficacy using a UAV autopilot system case study deployed on a quad-core platform (Raspberry Pi). Results show that the proposed scheme introduces minimal recovery overhead, measured in microseconds on a Raspberry Pi, and achieves 100% coverage in reconfiguration responses to compromised tasks in synthetic test cases.

CCS Concepts: • **Computer systems organization** → **Real-time systems**.

Additional Key Words and Phrases: Schedule Reconfiguration, Multicore, Security

ACM Reference Format:

Zain A. H. Hammadeh, Monowar Hasan, and Mohammad Hamad. 2025. RESCUE: A Reconfigurable Scheduling Framework for Securing Multi-Core Real-Time Systems. *ACM Trans. Cyber-Phys. Syst.* 1, 1, Article 1 (April 2025), 23 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

Real-time systems are essential to various critical applications, including automotive, aerospace, smart grids, industrial control, and space exploration. To ensure safe operation, real-time systems must meet strict timing requirements, commonly referred to as *deadlines*. However, due to their reliance on off-the-shelf components, increased connectivity, and high value to adversaries, these

This research is an extended version of the work published in the proceedings of ISORC 2024 [1]. The major improvement from the previous paper is the enhancement of the partitioning algorithm to be deadline-miss-tolerant for non-safety-critical tasks. The relaxation of the sufficient schedulability condition proposed in this work (see §3.4) significantly improves the coverage of RESCUE and improves the graceful degradation of the reconfiguration mechanism. To evaluate our proposed ideas, we added a new set of experiments (§4.1). Our results are reported in new figures (see Figure 4, Figure 6, Figure 7, Figure 9, and Figure 10). We also revised related work (§6) with the most up-to-date literature and made several revisions throughout the paper to improve clarity and readability.

Authors' addresses: Zain A. H. Hammadeh, zain.hajhammadeh@dlr.de, German Aerospace Center (DLR), Braunschweig, Germany; Monowar Hasan, monowar.hasan@wsu.edu, Washington State University, Pullman, USA; Mohammad Hamad, mohammad.hamad@tum.de, Technical University of Munich, Munich, Germany.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2025 ACM.

ACM 2378-962X/2025/4-ART1

<https://doi.org/XXXXXXXX.XXXXXXX>

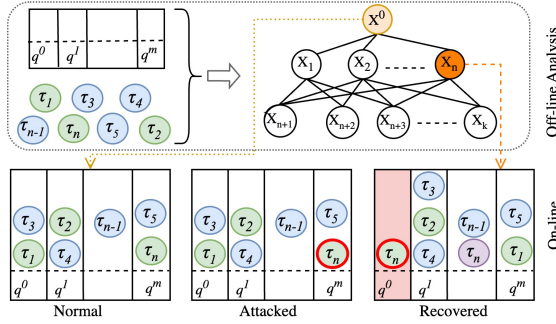


Fig. 1. Our proposed reconfiguration process (RESCUE) — offline analysis evaluates safety-critical (green) and non-safety-critical (blue) tasks and core counts, generating reconfiguration plans: a base schedule (χ^0) and compromised task responses. χ^0 remains in use until a task compromise, e.g., τ_n , triggers core q^0 isolation and activates recovery plan χ_n .

systems are becoming increasingly vulnerable to cyber-attacks. This vulnerability has been demonstrated by recent cyber-attacks targeting automobiles [2], drones [3], and medical robots [4].

A recent trend in real-time system design is the adoption of *multi-core* chips to improve performance and energy efficiency. Typically, safety-critical and non-safety-critical tasks coexist on a multi-core platform [5]. However, sharing resources between critical and non-critical tasks introduces security risks [6]. As demonstrated by researchers [2], compromising a non-critical task could potentially allow attackers to take control of the entire system [7]

Safety-critical applications require autonomy to withstand security attacks and operate in a “fail-safe” manner. Existing response options to thwart security breaches include run-time reconfiguration [8], task redundancy deployment [9], or task restart procedures [10, 11]. However, run-time reconfiguration lacks determinism — a critical requirement for real-time systems. Besides, implementing task redundancy can be prohibitively expensive and increases task response times, leading to missed deadlines. Moreover, restarting a malicious task (or the entire platform) does not eliminate the underlying vulnerability, as an adversary can repeat the attack. Furthermore, they are platform-dependent, require architectural modifications, and increase system downtime.

Due to the limitations of existing solutions, this paper investigates the following problem:

How can a real-time system running on a multi-core platform remain safe and operational (i.e., meet temporal constraints of all tasks) under security breaches where attackers compromise one or more tasks?

In response to the above problem, we propose a *reconfigurable scheduling* approach (named RESCUE) that operates at the software (viz., scheduler) level (i.e., does not require custom hardware). Our proposed technique *autonomously adjusts task schedules and core allocation* based on pre-computed “recovery plans.” Our idea stems from the fact that rather than immediately restarting a compromised task, it may be more beneficial to accumulate comprehensive *insights* into the attacker’s motives and the root cause of the attack. At the same time, the system is functional with the same level of autonomy (viz., does not miss any deadlines). This entails monitoring the compromised task to gain valuable information for devising effective attack mitigation strategies [12]. Hence, our technique will aid in systems *forensics*. Task termination becomes imperative at a certain stage due to escalated risks, a predefined threshold being reached, or sufficient information

being acquired. However, one concern is ensuring the attack's repercussions do not extend to other tasks allocated to the same resource (i.e., processor cores). Thus, we need a *recovery plan*. This plan outlines the steps to reschedule these tasks onto alternative resources, ensuring the system's continued functionality and security.

Contributions. We propose an offline analysis technique (RESCUE) that (a) explores the space of all feasible scheduling for a given set of real-time autonomous tasks running on a multi-core platform, and (b) computes the necessary recovery plans invoked during a suspicion of an attack. Each plan (we refer to it *configuration*) represents a feasible schedule of at least the safety-critical tasks while isolating compromised tasks on other cores. Figure 1 shows a case when the safety-critical task τ_n is got compromised. The system will be autonomously rescheduled to the recovery plan χ_n in which τ_n is isolated on core q^0 , and the other tasks, including a fresh instant of τ_n are scheduled on the other three cores.

In this paper, we made the following contributions.

- An offline analysis that explores the space of feasible scheduling for different compromised tasks (§3).
- A graceful degradation reconfiguration and security-aware deadline-miss-tolerant partitioning algorithm that allows the designers to integrate security constraints with guaranteeing maximum number of consecutive deadline misses. (§3.3).
- We performed a comprehensive evaluation using synthetic workloads and an autonomous system case study (UAV autopilot systems running on Raspberry Pi). Our approach enables run-time reconfiguration with microsecond-level overhead while covering up to 100% of potential responses for compromised tasks (§4).

2 MODEL AND ASSUMPTIONS

2.1 System Model

We consider a real-time system equipped with M identical cores $Q = \{q^0, q^1, \dots, q^{M-1}\}$ running a finite set of N independent real-time tasks \mathcal{T} under a real-time operating system (RTOS). Tasks are periodic and each task τ_i is defined as a tuple $\{C_i, D_i, T_i, \pi_i, \alpha_i\}$, where C_i is the worst-case execution time (WCET); D_i is the relative deadline; T_i is the period, such that $C_i \leq D_i \leq T_i$; π_i is the priority; α_i is the affinity set, which contains the indices of the cores on which τ_i can run. The affinity set can change during run-time. We consider the task level fixed-priority preemptive scheduling, i.e., each task τ_i is assigned a fixed priority π_i that does not change at run-time. The task set \mathcal{T} is sorted by task priorities, i.e., $\forall i : \pi_i \leq \pi_{i+1}$ where 0 is the highest priority.

The worst-case response time R_i^+ of a task τ_i is then the maximum response time among all jobs of τ_i . We consider a constrained deadline model, i.e., $D_i \leq T_i$. A task τ_i misses its deadline if and only if: $\exists \ell \in \mathbb{N}^+ : R_i^\ell > D_i$, where R_i^ℓ is the response time of the job ℓ . All tasks $\in \mathcal{T}$ can be assigned to any core. However, the worst-case execution time of τ_i does not change.

The tasks are classified into safety-critical and non-safety-critical tasks. We denote by τ^s (resp. $\tau^{\bar{s}}$) the safety-critical (resp. non-safety-critical) task. Consequently, we denote by \mathcal{T}^s (resp. $\mathcal{T}^{\bar{s}}$) the set of all safety-critical (resp. non-safety-critical) tasks, where

$$\mathcal{T} = \mathcal{T}^s \cup \mathcal{T}^{\bar{s}} : \mathcal{T}^s = \{\tau_i^s | \forall \tau_i \in \mathcal{T}\}, \mathcal{T}^{\bar{s}} = \{\tau_i^{\bar{s}} | \forall \tau_i \in \mathcal{T}\}.$$

We assign the tasks to the cores using partitioned scheduling with a task-level fixed priority scheduling policy. Hence, $\forall \tau_i \in \mathcal{T} : |\alpha_i| = 1$. Our partitioning algorithm is implemented as Satisfiability Modulo Theories (SMT) constraints (§3.3). The system may have *security constraints* (SC) that the partitioning algorithm has to consider. For example, a security constraint may state that two tasks τ_i and τ_j with different criticality must not be allocated on the same core. The system

may have any number of security constraints. The partitioned scheduling is backed by a *sufficient schedulability condition* to guarantee the hard real-time constraints of the safety-critical tasks. A non-safety-critical task is scheduled such that every at most μ_i consecutive deadline misses are followed by a deadline hit.

We assume that the tasks are subject to attacks. If a task got compromised, denoted $\hat{\tau}$, it will be isolated from other tasks. On the one hand, all non-safety-critical compromised tasks will be assigned to only one core q^ρ . Formally, if $\hat{\tau}_i^s$ and $\hat{\tau}_j^s$ then $\alpha_i \cup \alpha_j = \{q^\rho\}$. In other words, we have a best-effort approach to isolate non-safety-critical tasks, such that we isolate only what can be scheduled to one core; other tasks are ignored. On the other hand, we isolate every compromised safety-critical task on one core alone. Formally: if $\hat{\tau}_i^s$ then: $\forall \tau_j \in \mathcal{T} : \alpha_i \cap \alpha_j = \phi$. The system has a *safe-mode*, in which only the safety-critical tasks run on some cores to guarantee a fail-operational system. Table 1 includes all the notations used in the paper.

Table 1. Key mathematical notations used in this work.

Notation	Description	Notation	Description
M	Number of cores	Q	Set of all cores
q^ρ	The ρ -th core in Q	N	Number of tasks
\mathcal{T}	The set of all tasks	τ_i	Task
C_i	Worst-case execution time of τ_i	D_i	Relative deadline of τ_i
T_i	Period of τ_i	π_i	Priority of τ_i
α_i	Affinity set of τ_i	$\hat{\tau}_i^s$	Compromised non-safety-critical task
τ_i^s	Safety-critical task	\mathcal{T}^s	Set of all safety-critical tasks
$\hat{\tau}_i^s$	Non-safety-critical task	$\mathcal{T}^{\bar{s}}$	Set of all non-safety-critical tasks
$\hat{\tau}$	Compromised task	$\hat{\tau}_i^s$	Compromised safety-critical task
χ_l	Configuration (see Def. 1)	C_l	Combination (see Def. 2)
χ^σ	Safe-mode configuration (see Def. 4)	\mathcal{L}	Length of the critical path (see Def. 5)
\mathcal{D}	Degradation of efficiency of RESCUE	BW_i	Level- i Busy-window
μ_i	Tolerable consecutive deadline misses		

2.2 Threat Model

We assume that an adversary may compromise one or several specific tasks. For instance, the attacker can compromise tasks through memory corruption attacks, such as illegal memory writes, buffer overflows, and control-flow hijacking. Once a task is compromised, the attacker can attempt to manipulate its execution or interfere with other tasks running on the same core. However, we do not consider attacks that exploit shared system resources among different cores, such as last-level cache side-channel attacks, bus snooping, or other side/covert-channel attacks. Since these attacks operate across cores, they require additional mitigation strategies beyond our proposed solution. It is important to note that we do not make explicit assumptions about the methods by which a task could become malicious. For example, an attacker might exploit existing system vulnerabilities or bugs or use social engineering tactics [13]. The specific means by which a task can be compromised is outside the scope of this paper. Further, for multi-vendor systems, a rogue vendor could inject malicious logic that may trigger at run-time [14]. Our focus here is to “isolate” anomalous tasks and ensure other, time-critical, benign tasks can meet their temporal requirements, and hence, the system remains *safe*. We assume that the RTOS provides mechanisms that support the isolation of various processes. We further assume the existence of a detection mechanism running on a protected space (say, within the kernel) that triggers the reconfiguration process when a task has been compromised.




Fig. 2. Offline and online analysis phases and the various steps used in RESCUE.

Table 2. A toy system configuration (seven tasks running on four cores).

Task	C_i	$D_i = T_i$	π_i	α_i	t_i^{out}	Type
τ_0	10	50	0	[1]	20	Critical
τ_1	10	50	1	[0]	20	Critical
τ_2	15	50	2	[0]	30	Critical
τ_3	40	200	3	[0]	80	NonCritical
τ_4	60	150	4	[1]	120	NonCritical
τ_5	110	1000	5	[2]	220	NonCritical
τ_6	65	400	6	[3]	130	NonCritical

3 RESCUE: SECURITY-COGNIZANT SCHEDULING

This paper proposes a reconfigurable scheduling technique, RESCUE, that *reacts* to security attacks on real-time tasks. Our goal is to reschedule the non-compromised tasks on a subset of cores to isolate the compromised task/s on the other cores while guaranteeing the real-time requirements of other tasks. The compromised task will be isolated for a predefined period.

Table 2 shows a real-time taskset running on a quad-core platform ($M = 4$). In the task set, τ_0 and τ_2 must not be scheduled on the same core as a design requirement. Consider that τ_3 got compromised. Letting τ_3 continue running on q^0 with the safety-critical tasks τ_1, τ_2 is dangerous. Our solution aims to *reconfigure* the system to a new configuration in which we re-schedule the non-compromised tasks on cores q^0, q^1, q^2 and isolate the compromised task τ_3 on q^3 . Such isolation may save more auditing logs on the activity of τ_3 for forensics analysis.¹ When τ_3 is isolated, it cannot read/write to I/O and not interfere with any non-compromised task.

We now present how to compute the configuration, the reconfiguration methodology, and the timing analysis. We start with a few definitions.

3.1 Configurations

Defining a feasible schedule for the real-time tasks under partitioned scheduling requires finding feasible partitions. We must allocate each task to one core where its schedulability is guaranteed. Note that there might be more than one feasible schedule for a task set and a multi-core platform.

DEFINITION 1. For a given task set and a multi-core platform, a configuration, denoted by χ_l , is an allocation of the tasks to the cores such that all the real-time tasks are schedulable.

Hence, there might be more than one configuration for a given task set and a multi-core platform. We propose an SMT-based partitioning algorithm to compute a configuration. The algorithm uses a “busy-window”-like analysis as SMT constraints (§3.3). The schedule of tasks to the cores where

¹Post isolation forensics, however, is not within the scope of this work.

no compromised tasks might be given or computed using our partitioning algorithm. We refer to this configuration as the *basic configuration* χ_0 .

DEFINITION 2. For a given task set, a combination, denoted by C_l , is a set of tasks where each task could be compromised or not, i.e., the compromised safety-critical task will be represented by its compromised instance.

DEFINITION 3. If τ_i is not compromised in C_l and τ_i is also not compromised in C_k , we say $C_k \geq C_l$.

For every combination, we need to compute a new configuration. There are 2^N combinations representing a lattice equipped by the partial order \geq . However, we cannot find a feasible schedule for every combination because we aim to isolate the compromised tasks on some cores. Also, a minimum set of cores is needed for the safety-critical tasks to keep them running. Hence, we can compute configurations for a subset of the lattice before moving to a *safe-mode configuration*.

DEFINITION 4. A safe-mode configuration χ^σ is a configuration where only the safety-critical tasks, \mathcal{T}^s , are scheduled.

For a combination C_l , if there is no feasible schedule, we assign the safe-mode configuration χ^σ . The configurations also can be ordered in a similar way to the combinations: $\chi_k \geq \chi_l$ if $C_k \geq C_l$.

RESCUE builds on three algorithms to compute and connect the configurations in a lattice form. Figure 2 shows the interactions between the algorithms. Alg. I computes the feasible configuration by utilizing the security-aware scheduling presented in Alg. II. Alg. III illustrates the procedure to connect the configurations to compose the final output of the offline phase of RESCUE.

Algorithm I (BUILDCONFIGURATIONS). We define the configuration χ_l as a data structure as Line 3 shows:

- *key* is a unique name given to χ_l ;
- *comb* is the associated combination C_l ;
- *schedule* is the affinity set of the tasks according to χ_l ;
- *previous* is a pointer to the previous configuration χ_k where $\chi_k \geq \chi_l$; and
- *next* is a pointer to the next configuration χ_m , where $\chi_l \geq \chi_m$.

Note that χ_l may have multiple configurations in its next and previous fields. We start with computing the basic configuration χ_0 (Lines 6-10), in which we try to balance the load on the cores. The *comb* is assigned an empty set because there are no compromised tasks. Next, we compute the safe-mode configuration χ^σ (Lines 11-13). For χ^σ , we follow Definition 4 and consider only the safety-critical tasks (\mathcal{T}^s) and try to schedule them on the minimum number of cores. The rest of the algorithm tries to compute a feasible configuration χ_l for every combination C_l (Lines 15-31).

For every combination C_l : First, we check if there is a feasible schedule of the non-compromised tasks (Line 16). We try to minimize the number of cores needed to schedule them to make room to isolate the compromised tasks. If (a) there is no feasible schedule or the schedule uses all the cores, and (b) there is no room to isolate any compromised task, we skip this combination (Lines 17-20), i.e., we assign the safe-mode configuration χ^σ to C_l . Second, we start with isolating as much as possible of the compromised safety-critical tasks (Lines 22-27). We isolate, i.e., schedule, each compromised safety-critical task $\hat{\tau}_i^s$ on one core. Third, we try to isolate (schedule) as much as possible of compromised non-safety-critical tasks on one core (Lines 28, 29). Finally, we call Alg. III to connect the computed configurations.

Algorithm II (SECURITYAWARESCHEM). This algorithm represents the partitioning approach and the schedulability test, which is called in five places in Alg. I. The partitioning approach and the schedulability test are security-aware and SMT-based solutions (§3.3). However, we present Alg. II here for better readability.

Algorithm I: BuildConfigurations

```

1 Input: taskset  $\mathcal{T}$ , cores  $Q$ ,  $SC$ 
2 Output: lattice of configurations
3 Declare struct  $\chi$  {key, comb, schedule, previous, next}
4  $lattice \leftarrow \{\}$ 
5  $M \leftarrow |Q|$ 

6  $sched \leftarrow SecurityAwareSched(\mathcal{T}, Q, SC, balance)$ 
7 if not  $sched$  then
8   return  $lattice$ 
9  $\chi_0 \leftarrow \chi(key=ConfigBasic, comb=\{\}, schedule = sched, previous = None)$ 
10  $lattice \leftarrow lattice \cup \chi^0$ 

11  $sched \leftarrow SecurityAwareSched(\mathcal{T}^s, Q, SC, best)$ 
12  $\chi^\sigma \leftarrow \chi(key=ConfigSafe, comb=\mathcal{T}, schedule = sched, next=\chi^0)$ 
13  $lattice \leftarrow lattice \cup \chi^\sigma$ 
14  $Combs \leftarrow computeAllCombs(\mathcal{T})$ 

15 for  $C_l \in Combs$  do
16    $sched, Q_a \leftarrow SecurityAwareSched(\mathcal{T}^s \cup \{\tau_i^s | \forall \tau_i \in C_l\}, Q, SC, best)$ 
17   if not  $sched$  then
18     continue
19   if  $Q_a = \emptyset$  then
20     continue
21    $sched^s \leftarrow \{\}$ 
22   for  $\hat{\tau}_i^s \in \{\hat{\tau}^s | \forall \hat{\tau}^s \in C\}$  do
23     while  $Q_a \neq \emptyset$  do
24        $q \leftarrow Q_a.pop$ 
25        $\hat{\alpha}_i^s \leftarrow q.index$ 
26        $sched^s \leftarrow \hat{\alpha}_i^s$ 
27     break
28   if  $Q_a \neq \emptyset$  then
29      $sched^{\hat{s}} \leftarrow SecurityAwareSched(\{\hat{\tau}^s | \forall \hat{\tau}^s \in C\}, Q_a, best)$ 
30    $\chi_l \leftarrow \chi(key=Config+C, comb=C, schedule = sched \cup sched^s \cup sched^{\hat{s}})$ 
31    $lattice \leftarrow lattice \cup \chi$ 

32  $lattice \leftarrow ConnectConfigurations(lattice)$ 
33 return  $lattice$ 

```

The algorithm can work in two modes: *balance* and *best*. In the *balance* mode, the algorithm calls the *call_smt* procedure to allocate the tasks into all cores (Lines 4, 5). This mode is used for the basic configuration χ_0 . In the *best* mode, the algorithm calls the *call_smt* procedure in a loop (Lines 6-13). In each iteration, one more core is added to the set of cores given to the SMT solver. It stops once the SMT solver can find a solution, which guarantees that the tasks are scheduled on the minimum number of cores. It stops with an empty schedule, i.e., it fails if there are no more cores exist.

Algorithm II: SecurityAwareSched

```

1 Input: taskset  $\mathcal{T}$ , cores  $Q$ ,  $SC$ , approach
2 Output: sched,  $Q_a$ 
3 New smt_solver
4 if approach = balance then
5   success, sched = call_smt( $\mathcal{T}$ ,  $Q$ ,  $SC$ )
6 if approach = best then
7   success ← False
8    $Q_u \leftarrow \{\}$ 
9   while not success do
10    if  $Q = \emptyset$  then
11      break
12     $Q_u \leftarrow Q_u \cup Q.pop()$ 
13    success, sched = call_smt( $\mathcal{T}$ ,  $Q_u$ ,  $SC$ )
14  $Q_a \leftarrow Q$ 
15 return sched,  $Q_a$ 
16
17 call_smt( $\mathcal{T}$ ,  $Q$ ,  $SC$ ):
18 for  $q^p \in Q$  do
19   for  $\tau_i \in \mathcal{T}$  do
20     smt_solver.add_constraint( $\sum_{j \leq i} \lceil \frac{D_i}{T_j} \rceil \cdot C_j \cdot x_j^p \leq D_i$ )
21 for  $sc \in SC$  do
22   smt_solver.add_constraint(sc)
23 success, sched = smt_solver.solve()
24 return success, sched

```

Lines 17-24 present the *call_smt* procedure. It builds the scheduling constraints (Lines 18-20) and the security constraints (Lines 21, 22). We call the SMT solver in Line 23.

Algorithm III (CONNECTCONFIGURATIONS). This is the last algorithm in the offline phase of RESCUE. It takes the configurations generated by Alg. I and connects them depending on the \geq relation. For each configuration χ_l , we search in the lattice for a parent configuration χ_k , i.e., $\chi_k \geq \chi_l$. We update the *previous* attribute of the child (χ_l) and the *next* of the parent (χ_k), where the *key* is the index of the compromised task (Lines 5-8). The configuration χ_l may have more than one parent depending on the number of compromised tasks in C_l . If the configuration has no parents, we delete the configuration (Lines 9, 10). For each configuration χ_l that does not have any child, it has to connect to the safe-mode configuration χ^σ , i.e., add χ^σ to $\chi_l.next$. Also, if the length of the *next* attribute of χ_l is smaller than the number of compromised tasks in C_l (Line 19), then there are unfeasible configurations that had to be children of χ_l . In this case, we add the χ^σ to replace the missing children. In other words, the unfeasible configurations are replaced by the safe-mode configuration.

In the example task set (Table 2), the schedule of tasks to cores represents the basic configuration χ_0 . In this example, there are 128 feasible configurations. If τ_3 is compromised, the system will move to the configuration χ_4 in which τ_3 is isolated on q^3 . According to the proposed algorithms (Alg. I - Alg. III) we have: $\chi_0.next[3] = \chi_4$ and $\chi_4.previous[3] = \chi_0$. The proposed algorithms generated

Algorithm III: ConnectConfigurations

```

1 Input: lattice
2 Output: lattice
3 for  $\chi_l \in \text{lattice}$  do
4   for  $\chi_k \in \text{lattice}$  do
5     if  $\chi_k \geq \chi_l$  then
6        $\text{key} = \chi_k.\text{comb} \setminus \chi_l.\text{comb}$ 
7        $\chi_l.\text{previous}[\text{key}] = \chi_k$ 
8        $\chi_k.\text{next}[\text{key}] = \chi_l$ 
9
10    if  $\chi_l \neq \chi_0$  &  $\chi_l \neq \chi^\sigma$  &  $!\chi_l.\text{previous}$  then
11      delete  $\chi_l$ 
12
13    if  $\chi_l \neq \chi^\sigma$  then
14      if  $!\chi_l.\text{next}$  then
15         $\chi_l.\text{next}[\text{all}] = \chi^\sigma$ 
16         $\chi^\sigma.\text{previous}[\text{all} + \text{counter}] = \chi_l$ 
17         $\text{counter} \leftarrow \text{counter} + 1$ 
18      else
19        if  $|\mathcal{T}| - |\{\hat{\tau} | \forall \hat{\tau} \in C\}| > |\chi_l.\text{next}|$  then
20           $\chi_l.\text{next}[\text{all}] = \chi^\sigma$ 
21           $\chi^\sigma.\text{previous}[\text{all} + \text{counter}] = \chi_l$ 
22           $\text{counter} \leftarrow \text{counter} + 1$ 
23
24  return lattice

```

the full connected lattice of configurations within 1.3 seconds. The size of the generated lattice is 4.696 kilobytes.

3.2 Reconfiguration Methodology

As explained earlier, the configurations, including the safe-mode configuration, are computed offline. We define a time-out for each task, denoted t_i^{out} , which designers can specify based on the application. The time-out period represents how long we keep the compromised tasks isolated before killing the job and keeping the new job of the safety-critical task or starting a fresh job of the non-safety-critical task. The reconfiguration is a procedure triggered when a task is compromised and after the time-out. Alg. IV presents the reconfiguration procedure. During the reconfiguration, we move from the configuration χ_k to χ_l after the task τ_i got compromised if $\chi_k \geq \chi_l$. Similarly, we move from the configuration χ_l to χ_k after t_i^{out} .

Algorithm IV (RECONFIGURATION). This algorithm represents the online phase of RESCUE and is integrated with the scheduler. The algorithm utilizes the lattice of configurations (generated by the offline phase) and an *event*. If the event is *Isolate*, i.e., τ_i got compromised, then the algorithm moves to the child configuration in the *next* attribute of the current configuration depending on the index of the compromised task i (Lines 5-11). After a timeout t_i^{out} , a new event will be triggered

Algorithm IV: Reconfiguration

```

1 Input: lattice, current config:  $\chi$ , event, taskID:  $i$ 
2 Output: schedule
3  $\chi_0 \leftarrow \text{lattice}[\text{ConfigBasic}]$ 
4  $\chi^\sigma \leftarrow \text{lattice}[\text{ConfigSafe}]$ 
                                     /* task got compromised */
5 if event is Isolate then
6   if  $\chi.\text{key}$  is ConfigSafe then
7     return  $\chi^\sigma$ .schedule
8   if  $i$  in  $\chi.\text{next}$  then
9     return  $\chi.\text{next}[i]$ .schedule
10  else
11    return  $\chi^\sigma$ .schedule
                                     /* time-out */
12 if event is Integrate then
13   if  $\chi.\text{key}$  is ConfigSafe then
14     return  $\chi_0$ .schedule
15   if  $i$  in  $\chi.\text{previous}$  then
16     return  $\chi.\text{previous}[i]$ .schedule

```

to *Integrate* τ_i . In this case, the algorithm moves to the parent configuration in the *previous* of the current configuration (Lines 12-16).

In our example (Table 2), if τ_3 got compromised, the schedule moves to the configuration χ_4 . After a timeout = $\tau_3^{\text{out}} = 80$, a new reconfiguration is triggered, in which the scheduler moves back to the configuration χ_0 .

Let M_{\min} denote the minimum number of cores needed to schedule all tasks; hence, $M_{\min} \leq M$. We bound the number of isolated tasks using the following lemma.

LEMMA 1. *The maximum number of compromised safety-critical tasks that RESCUE can isolate is $M - M_{\min}$.*

PROOF. We isolate each compromised safety-critical task $\hat{\tau}^s$ on one core, and we need at least M_{\min} cores to schedule the task set. If some non-safety-critical tasks got compromised, we give the priority to isolate the safety-critical tasks first. \square

LEMMA 2. *RESCUE can cover all 2^N combinations and isolate all compromised tasks if $1 + |\mathcal{T}^S| + M_{\min} \leq M$ and the compromised non-safety-critical tasks can fit in one core.*

PROOF. If such an over-provisioned system exists, the proposed approach could cover all combinations and isolate all compromised task. \square

An interesting indicator of the resilience of the schedule reconfigurability against security attacks is the length of the *critical path*, denoted by \mathcal{L} , as we formally define below.

DEFINITION 5. *The critical path is a sequence of configurations starting from the basic configuration with no compromised tasks and ending in the safe-mode.*

The length of the critical path indicates the number of compromised tasks the system can tolerate before entering the safe-mode. This length, \mathcal{L} , depends on the number of cores M , the

schedulability of the task set, and the security constraints. For systems not equipped with proper reaction mechanisms like RESCUE, the best case scenario is that the system will move to the safe-mode after the first attack. Hence, the length of the critical path is equal to one. We use the maximum number of compromised non-safety-critical tasks that may be excluded after a reconfiguration, i.e., neither executed nor isolated, as an indicator of the degradation of the efficiency of RESCUE, denoted by \mathcal{D} . In the experiments, we use these two indicators, \mathcal{L} and \mathcal{D} to study the efficiency of RESCUE.

3.3 Partitioning Algorithm

The busy-window [15] under the fixed priority preemptive (FPP) scheduling is computed as follows:

$$BW_i^{n+1} = \sum_{j \leq i} \left\lceil \frac{BW_i^n}{T_j} \right\rceil \times C_j, \quad (1)$$

with $BW_i^0 = C_i$. The recurrence stops at the convergence (i.e., $BW_i^{n+1} = BW_i^n$) or when $BW_i > D_i$ (i.e., the task τ_i is unschedulable).

LEMMA 3. For constrained-deadline periodic tasks, τ_i meets its deadline in the worst-case if:

$$BW_i \leq D_i. \quad (2)$$

PROOF. In FPP, if $BW_i \leq T_i$ then BW_i has one job and $BW_i = R_i^+$. Hence, if Eq. (2) holds, schedulability is guaranteed. \square

LEMMA 4. Let us consider a fixed-size window of D_i . For constrained-deadline periodic tasks, τ_i meets its deadline if:

$$\sum_{j \leq i} \left\lceil \frac{D_i}{T_j} \right\rceil \times C_j \leq D_i. \quad (3)$$

PROOF. For constrained-deadline periodic tasks, if the accumulative load within a window of D_i can finish before D_i , then τ_i meets its deadline, i.e., the condition in Eq. (2) is satisfied. \square

The sufficient condition in Eq. (3) does not need iteration to be computed. However, it is more pessimistic than the sufficient condition in Eq. (2), because, if τ_i meets its deadline, i.e., Eq. (2) is satisfied:

$$BW_i \leq \sum_{j \leq i} \left\lceil \frac{D_i}{T_j} \right\rceil \times C_j. \quad (4)$$

3.3.1 *SMT constraints.* Now, we have a linear non-iterative sufficient schedulability condition for τ_i . To map τ_i to the core q^p , we have to guarantee that Eq. (3) is satisfied for all higher priority tasks already mapped to q^p . We can formulate the sufficient condition in Eq. (3) as a linear constraint as follows:

$$\sum_{j \leq i} \left\lceil \frac{D_i}{T_j} \right\rceil \times C_j \times x_j^p \leq D_i, \quad (5)$$

where $x_j^p \in \{0, 1\}$ a binary variable indicates whether τ_j is mapped to the core q^p or not. Hence, we can define the partitioned scheduling for periodic tasks under FPP as an SMT with $M \times N$ variables and constraints for M cores and N tasks.

The SMT-based partitioned scheduling has the following constraints:

$$\forall q^p \in Q \forall \tau_i \in \mathcal{T} : \sum_{j \leq i} \left\lceil \frac{D_i}{T_j} \right\rceil \times C_j \times x_j^p \leq D_i. \quad (6)$$

Also, the SMT integrates security constraints, e.g., τ_i and τ_j are not allowed to run on the same core as follows:

$$\neg(x_j^p \wedge x_i^p). \quad (7)$$

3.4 Graceful Degradation

As stated in our system model (see §2.1), non-safety-critical tasks can tolerate few deadline misses, namely μ_i consecutive deadline misses, as long as they are followed by at least one deadline hit. Therefore, non-safety-critical tasks follow the *weakly-hard real-time model* [16]. The partitioning algorithm presented so far is pessimistic in scheduling the non-safety-critical tasks, as it considers them all as hard real-time tasks. Besides, RESCUE sacrifices one or more non-safety-critical tasks to schedule safety-critical tasks. The conservative scheduling of the partitioning algorithm limits the feasible configurations, which pushes RESCUE to sacrifice non-safety-critical tasks. However, as we present next, it is possible to provide some graceful degradation of RESCUE by relaxing the sufficient scheduling condition.

To achieve graceful degradation, we update the partitioning algorithm such that it schedules a non-safety-critical task as long as it has at least one deadline hit after at most μ_i consecutive deadline misses. In the weakly-hard real-time model, the task is allowed to miss at most μ_i deadlines in every window of K consecutive executions. In our case, the non-safety-critical task has the weakly-hard constraint $(\mu_i, \mu_i + 1)$. Such tasks are known in the literature as *high-tolerance weakly-hard tasks* [17] because $\mu_i / \mu_i + 1 \geq 0.5$. To update the constraints in the SMT-based scheduling, we consider henceforth that each task τ_i has a parameter μ_i that represents the maximum allowed number of consecutive deadline misses. For all safety critical tasks $\mu_i = 0$. For all non-safety-critical tasks $\mu_i \geq 0$.

LEMMA 5. *For deadline-constrained periodic tasks, τ_i misses only at most μ_i deadlines in a row if:*

$$BW_i \leq (\mu_i + 1)D_i \quad (8)$$

PROOF. For deadline-constrained tasks, BW_i extends to include the next job if and only if the previous jobs missed their deadlines. When $BW_i \leq (\mu_i + 1)D_i$, then BW_i contains at most $\mu_i + 1$ jobs. Also, $\forall \geq 0$ the last job meets its deadline because the response time of the $(\mu_i + 1)$ -th job in BW_i is equal to $BW_i - \mu_i \times D_i$. Hence, there are at most μ_i consecutive deadline misses followed by one deadline hit. \square

As BW_i is computed as a fixed-point equation, the following condition is satisfied for deadline-constrained periodic tasks.

$$BW_i \leq \sum_{j \leq i} \left\lceil \frac{(\mu_i + 1)D_i}{T_j} \right\rceil \times C_j. \quad (9)$$

LEMMA 6. *For deadline-constrained periodic tasks, τ_i misses only at most μ_i deadlines in a row if:*

$$\sum_{j \leq i} \left\lceil \frac{(\mu_i + 1)D_i}{T_j} \right\rceil \times C_j \leq (\mu_i + 1)D_i \quad (10)$$

PROOF. Similar to Lemma 4. \square

The SMT-based partitioned scheduling considering missing deadlines has the following constraints:

$$\forall q^p \in Q \forall \tau_i \in \mathcal{T} : \sum_{j \leq i} \left\lceil \frac{(\mu_i + 1)D_i}{T_j} \right\rceil \times C_j \times x_j^p \leq (\mu_i + 1)D_i \quad (11)$$

Note that the new set of constraints is more general than the constraints in Eq. (6) as we can get Eq. (6) from Eq. (11) by substituting μ_i by zero for all non-safety-critical tasks.




Fig. 3. The schedulability ratio of RESCUE vs the APA scheduling [19]. RESCUE outperforms the APA scheduling for high utilization. The schedulability ratio of RESCUE and APA scheduling decreases when U/N increases.




Fig. 4. The schedulability ratio of RESCUE considering tolerable deadline misses. The schedulability ratio of RESCUE decreases when U/N increases.

4 EXPERIMENTAL EVALUATION

We evaluate RESCUE on two fronts: (a) synthetically generated workload for broader design-space exploration (§4.1) and (b) case study on a real-time platform (§4.2).

4.1 Synthetic Test Cases

Workload Generation. We randomly generated synthetic test cases to carry out our experiments as follows: we distributed the utilization over the N tasks using the UUniFast algorithm [18]. Periods were generated randomly as follows: $T_i = f \times p$ where $f \in \{1, 2, 3, 4, 5, 6, 7, 8\}$, and $p \in \{280, 340, 450, 500\}$. Therefore, the periods of the task set or a subset of it can be harmonic. Relative deadlines were implicit $D_i = T_i$. Priorities were assigned according to the deadline monotonic (DM) approach. The number of critical tasks was decided randomly so that it is at most 40% of the total tasks.

4.1.1 Schedulability of the Partitioning Algorithm. The proposed security-aware SMT-based partitioning scheduling utilizes a linear sufficient condition of schedulability using Eq. (3), which is more pessimistic than, e.g., the original busy-window analysis, i.e., Eq. (2). This experiment tests the schedulability of our partitioning algorithm compared to the Arbitrary Processor Affinity (APA) scheduling [19]. APA scheduling dominates partitioned scheduling. However, the schedulability test of APA scheduling, which depends on the response time analysis of global scheduling [20], is very pessimistic. Hence, it may report a lower schedulability ratio than the partitioned scheduling. In

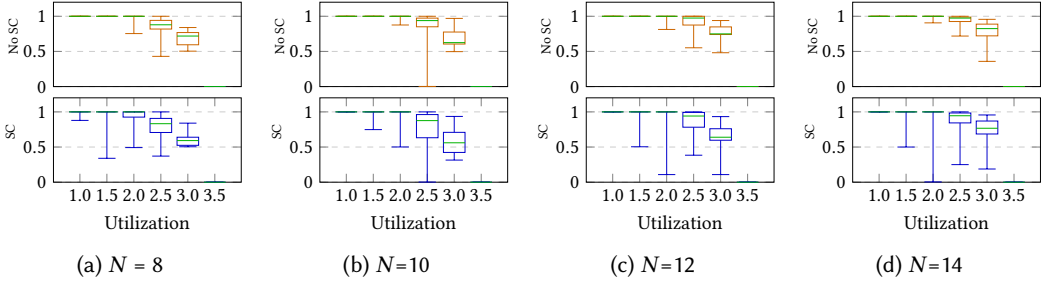


Fig. 5. Coverage ratio of RESCUE for $M = 4$. The coverage ratio represents the ratio between the number of feasible configurations, i.e., the size of the lattice, over 2^N . The green line represents the median. RESCUE can achieve coverage ratio of 100% for low utilization. The coverage ratio is related to the schedulability ratio.

this experiment, we support our decision to use SMT-based partitioned scheduling by showing the schedulability ratio of the APA scheduling, namely the heuristic-based scheduling, presented in [19], and RESCUE. We ran the experiments for $M = \{4, 6, 8\}$. Figure 3 shows the results. We generated 100 synthetic test cases for each value of the utilization between $M/2 \leq U \leq M$ with step=0.5. Figure 3 reports the results for $N = \{16, 20, 24, 28\}$. APA scheduling cannot outperform RESCUE. In contrast, the SMT-based partitioned scheduling in RESCUE can outperform the APA scheduling, e.g., for $U = 3.0$ when $M = 4$. The schedulability ratio of RESCUE and APA scheduling decreases when U/N increases. Figure 4 shows the improvement on the schedulability when considering that non-safety-critical tasks can tolerate deadline misses, i.e., $\mu > 0$. The results show the improvement in schedulability when the partitioning algorithm relaxes the schedulability constraints on the non-safety-critical tasks by allowing no more than μ_i consecutive deadline misses. Another notable remark w.r.t. APA is that APA is not extendable to consider tolerable deadline misses the way we did with the SMT-based scheduling. In fact, the majority of weakly-hard scheduling analysis techniques are designed for fixed-priority preemptive scheduling, and only a few approaches exist for dynamic schedulers such as earliest deadline first (EDF) [21].

4.1.2 Scalability and Efficiency of RESCUE. We generated 100 test cases for each value of the utilization between $M/2 \leq U \leq M$ with step=0.5. We repeated the experiment for $N = \{8, 10, 12, 14\}$ tasks and for $M = \{4, 6, 8\}$. As a security constraint, we considered three safety-critical tasks to run on three different cores. Figure 5 presents the coverage ratio for $M = 4$. The plot boxes in orange (the upper figures) consider no security constraints, while the blue plot boxes report the results considering the security constraints. For all tasks, RESCUE can achieve a coverage ratio of 100% for low utilization. However, the coverage ratio drops with the high utilization values. The coverage ratio is explained by the schedulability ratio of the security-aware SMT-based partitioning algorithm. In addition, security constraints impact the coverage ratio because they impose limitations on the scheduling options.

In Figure 6 and Figure 7, we repeated the experiments considering $\mu = 1$ and $\mu = 2$ resp. The coverage ratio increases for both cases, with no security constraints (in orange) and with security constraints (blue). The more the partitioning algorithm is tolerant of deadline misses, the higher the coverage ratio. We can observe that the coverage ratio for $\mu_i = 2$ (Figure 7) is higher than the coverage ratio for $\mu_i = 1$ (Figure 6). This improvement is explainable by the better schedulability ratio that the partitioning algorithm can achieve when it allows the non-safety critical tasks to have bounded consecutive deadline misses, see Figure 4.

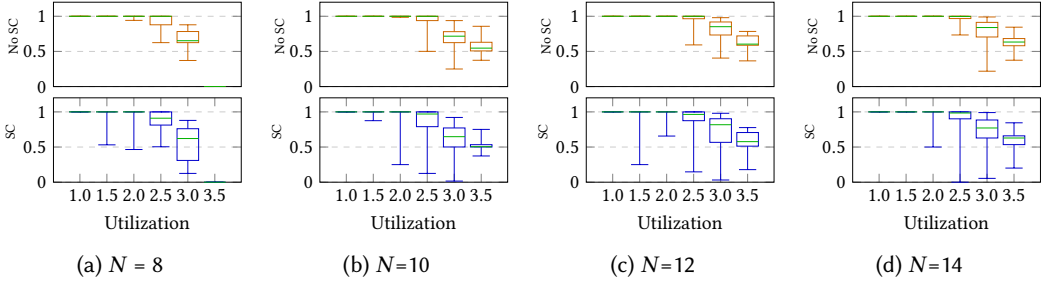


Fig. 6. Coverage ratio of RESCUE for $M = 4$ and $\mu = 1$. In this experiment, every non-safety-critical task will be guaranteed to have at least one deadline hit after each deadline miss in the worst-case.

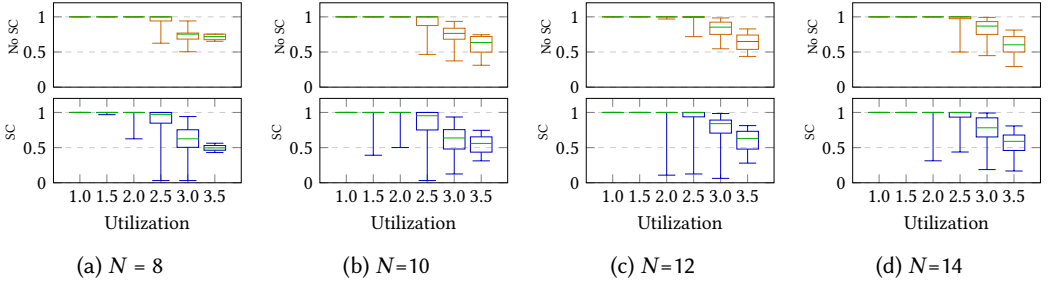




Fig. 7. Coverage ratio of RESCUE for $M = 4$ and $\mu = 2$. In this experiment, every non-safety-critical task will be guaranteed to have at least one deadline hit after 2 consecutive deadline misses in the worst-case.



(a) The length of the critical path \mathcal{L} . The base line represents the value of \mathcal{L} for systems without RESCUE. (b) The degradation of the efficiency \mathcal{D} . In the worst-case, \mathcal{D} can be equal to the number of non-safety-critical tasks.

Fig. 8. \mathcal{L} and \mathcal{D} for $M = 4$. For instance, $U = 1.0$, the values are reported in order from left to right for $N = \{8, 10, 12, 14\}$. The green line represents the median.

We studied the length of the critical path \mathcal{L} as an indicator of the resilience of RESCUE. Figure 8a shows \mathcal{L} for $M = 4$. For instance, $U = 1.0$, the values are reported in order from left to right for $N = \{8, 10, 12, 14\}$. The baseline represents the value of \mathcal{L} for systems that are not equipped with proper reaction mechanisms such as RESCUE. When the utilization is low, RESCUE can react to



(a) The length of the critical path \mathcal{L} . The base line represents the value of \mathcal{L} for systems without RESCUE.

(b) The degradation of the efficiency \mathcal{D} . In the worst-case, \mathcal{D} can be equal to the number of non-safety-critical tasks.

Fig. 9. \mathcal{L} and \mathcal{D} for $M = 4$ and $\mu = 1$. For instance, $U = 1.0$, the values are reported in order from left to right for $N = \{8, 10, 12, 14\}$. The green line represents the median.


any scenario of compromising tasks. However, for large values of utilization like $U = 3.0$, RESCUE does not have many options than moving to the safe-mode.

Figure 8b shows the degradation of the efficiency of RESCUE in the form of the number of compromised non-critical tasks that are neither scheduled nor isolated. In the worst-case scenario, \mathcal{D} can be equal to the number of non-safety-critical tasks. This happens when RESCUE moves to the safe-mode. However, the upper whisker ≤ 1 , which means that for 75% of the test cases RESCUE misses no more than one task when configuring the system after an attack.

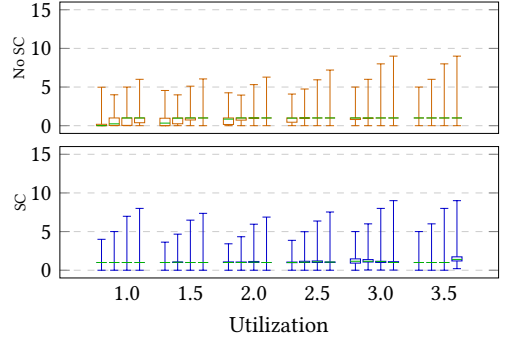
Figure 9 and Figure 10 report \mathcal{L} and \mathcal{D} considering $\mu_i = 1$ and $\mu_i = 2$ respectively. The length of the critical path increases with increasing μ_i . Relaxing the constraints on scheduling the non-safety-critical tasks by μ_i consecutive deadline misses extends the critical paths, which enhances the resilience of RESCUE. \mathcal{D} does not show significant change compared to the results reported in Figure 8b, i.e., $\mu_i = 0$. The relaxation in the schedulability constraints presented in Section 3.4 cannot illuminate or control the number of compromised non-critical tasks that are neither scheduled nor isolated. Hence, \mathcal{D} cannot reflect the improvement on the efficiency of RESCUE when $\mu_i > 0$.

To study the scalability of RESCUE, we computed the execution time of the offline phase for $N = \{8, 10, 12, 14\}$ tasks and $M = \{4, 6, 8\}$. The experiment ran on a general-purpose Linux server (18-core Intel Xeon CPU and 251 GB Memory). Figure 11 shows the execution time for $M = 4$. The sub-figures show a similar trend. Hence, we explain the results in Figure 11b. The median in the orange plot boxes (no security constraints) increases with the utilization. This can be explained by Alg. II (Lines 8-13) because increasing the utilization requires more cores to schedule the task set, which imposes more iterations on the *while*-loop. The decreasing median for the blue boxes, where the security constraints are considered, is mainly due to the low coverage ratio, i.e., the small size of the lattice (see Figure 5). Also, The execution time of the SMT solver scales with the number of constraints. Hence, the SMT solver consumes more time to find a feasible partitioning for the tasks when there are security constraints.

Increasing the number of cores for a given number of tasks increases the number of iterations that the *while*-loop takes in Alg. II (Lines 8-13). The main factor affecting the scalability is the number of tasks N , where the execution time increases with N . We now discuss the implementation details of RESCUE to explain this exponential growth.



(a) The length of the critical path \mathcal{L} . The base line represents the value of \mathcal{L} for systems without RESCUE.



(b) The degradation of the efficiency \mathcal{D} . In the worst-case, \mathcal{D} can be equal to the number of non-safety-critical tasks.

Fig. 10. \mathcal{L} and \mathcal{D} for $M = 4$ and $\mu = 2$. For instance, $U = 1.0$, the values are reported in order from left to right for $N = \{8, 10, 12, 14\}$. The green line represents the median.

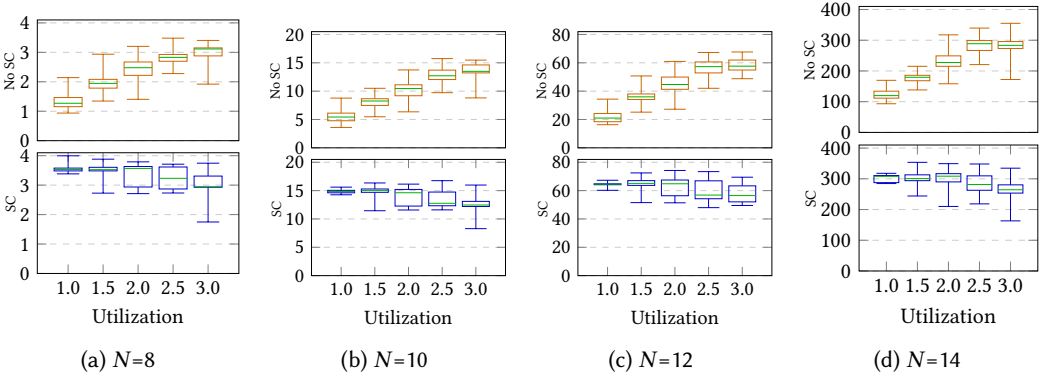


Fig. 11. Execution time of the offline phase in RESCUE for $M = 4$. The green line represents the median. The main factor is the number of tasks N . Also, number of constraints in the SMT solver, and the utilization are the main players.

RESCUE simulation engine is implemented in Python. The Algorithm BUILDCONFIGURATIONS (Alg. I) is implemented such that the for-loop works in parallel, utilizing all available cores. However, the Algorithm CONNECTNODES (Alg. III) is the bottleneck in RESCUE. Results show that Alg. III consumes no more than 2% of the execution time for $N=8$, and up to 43% for $N=14$. For $N=18$, the percentage can reach 92%. Adding parallelism to Alg. III is left for future development. However, the execution time of the offline phase has no impact on the online phase. Figure 12 depicts the lattice size as the output of Alg. III. The reported numbers represent the maximum size of the lattice for several tasks. The experiments show that the size of the lattice has no relation to the number of cores. The size of the lattice grows exponentially with the number of tasks. However, the size of the file is acceptable for nowadays embedded systems.

Note. Although RESCUE’s offline analysis scales exponentially with the number of tasks, we stress that the calculation occurs during the design phase (before the system deployment). Hence,




Fig. 12. The maximum lattice size in kB.

Table 3. Task parameters for the ArduCopter UAV case study.

Task	C	$T = D$	Type
rc_loop	130	4000	Critical
throttle_loop	75	20000	Critical
update_GPS	200	20000	Critical
update_optical_flow	160	5000	Critical
update_altitude	140	100000	Critical
run_nav_updates	100	20000	Critical
update_thr_average	90	10000	NonCritical
three_hz_loop	75	333333	NonCritical
compass_accumulate	100	5000	NonCritical
barometer_accumulate	90	20000	NonCritical
update_notify	90	5000	NonCritical
ekf_check	75	100000	NonCritical
landinggear_update	75	100000	NonCritical
lost_vehicle_check	50	100000	NonCritical
gcs_check_input	180	2500	NonCritical
gcs_send_heartbeat	110	5000	NonCritical

the runtime behavior is not impacted, as the time needed to adopt and perform the reconfiguration is independent of the number of tasks and can be accomplished within a few microseconds, even for large systems.

4.2 Case Study

We use the ArduCopter² UAV autopilot system as a case study. Table 3 shows the tasks considered in this case study. The tasks are periodic with implicit deadlines. As vanilla ArduCopter does not distinguish critical-vs-non-critical tasks, we manually inspected source code and documentation and classified them into critical and non-critical tasks based on their functionalities for our demonstration purposes. Also, we consider a quad-core platform. In this case study, the tasks are considered independent and implemented as threads.

We built a lattice of 65536 feasible configurations with 100% coverage ratio. Building the lattice took 1429.26 s. The critical path from the basic configuration to the safe-mode consists of 16

²<https://github.com/ArduPilot/ardupilot/blob/master/ArduCopter/Copter.cpp>




Fig. 13. The overhead of the reconfiguration procedure as measured on a multi-core embedded platform (Raspberry Pi).

configurations, i.e., $\mathcal{L} = 16$, including the basic and the safe-mode. Hence, compromising any combination of safety-critical and non-safety-critical tasks will lead to a feasible configuration. The system only goes into safe-mode when the attack can compromise all tasks. The lattice size (as a file) is 2.6 MiB.

4.2.1 Attack Scenario and Findings. We consider that the task `run_nav_update` depends on the task `update_GPS` and the task `update_optical_flow`, such that compromising `update_GPS` or `update_optical_flow` will help the attacker to compromise the task `run_nav_update`, especially if they are running on the same core, i.e., sharing the same L1 cache.

Let us consider the tasks `update_optical_flow` and `run_nav_update` are compromised after t_{prop} time. If the system is not equipped with the reconfiguration mechanism, it will be in the safe-mode of the system. However, using our technique, the system will reconfigure to a feasible schedule where `update_optical_flow` and `run_nav_update` are isolated on two different cores, i.e., two new (viz., uncompromised) instances of these two safety-critical tasks will be scheduled. In particular, the following reconfigurations will take place:

- Reconfigure from the basic configuration *ConfigBasic* to *Config_1_3* when the task `update_optical_flow` got compromised; let this point of time be t_0 .
- Reconfigure from *Config_1_3* to *Config_2_62* when the task `run_nav_update` got compromised; let this point of time be $t_1 > t_0$.
- After a period of time $t_0 + t_{update_optical_flow}^{out}$, the system moves to the configuration *Config_1_9*, in which only `run_nav_update` is isolated.
- After a period of time $t_1 + t_{run_nav_update}^{out}$, the system moves to *ConfigBasic*.

As we know the critical relation between the task `update_optical_flow` and `run_nav_update`, we will add a constraint on the SMT solver such that these two tasks are not allocated to the same core. We add a second constraint for `update_GPS` and `run_nav_update`. The new lattice of configurations also $\mathcal{L} = 16$. Hence, we can reduce the probability of compromising `run_nav_update`.

4.2.2 Reconfiguration Overhead. To measure the reconfiguration overhead, we tested the reconfiguration procedure (Alg. IV) on an embedded system (Raspberry Pi3) with quad-core Cortex-A72 (ARM v8) 64-bit SoC, working at a speed of 1.8 GHz. We repeated the test 3000 times and computed the execution time. Figure 13 shows the result. The overhead is negligible as it is at the microsecond level, which could be acceptable for many real-time autonomous applications.

5 DISCUSSION

RESCUE works in coordination with an intrusion detection system (IDS) [22]. This is a common assumption for many scheduler-based real-time security research [23–28]. We assume the presence of an IDS capable of detecting attacks and triggering RESCUE. Implementing such an IDS is beyond the scope of this paper. We note that the effectiveness of the system depends on the quality of the IDS. If the IDS has a high false positive rate, RESCUE may be triggered by incorrect detections.

Another important issue to consider is the predictability of reconfiguration in RESCUE. If an attacker gains knowledge of the reconfiguration plans, they could repeatedly compromise tasks on the same core to trigger frequent reconfigurations, potentially leading to a denial-of-service attack. However, as demonstrated in this paper, we have analyzed the number of steps (i.e., the number of tasks that need to be compromised) required to reach such a critical state. Our results show that our framework significantly enhances security compared to a system without it (see Figure 8). Additionally, RESCUE isolates compromised tasks and logs them for future updates. However, adaptive attacks are not entirely prevented. To further mitigate this issue, one approach is to keep the reconfiguration plans secure. While an adaptive attacker may eventually infer patterns over time, increasing the unpredictability of the scheduling process can further strengthen security. This can be achieved by generating multiple reconfiguration plans and selecting them dynamically at runtime. However, this approach introduces additional design-time overhead for computing alternative plans and may require extra storage during execution. Integration of such a scheme and analysis of security-performance trade-offs requires further research.

6 RELATED WORK

RESCUE is first introduced in our early work [1]. This paper extends RESCUE for cases where some non-safety-critical tasks can miss deadlines (albeit in a bounded manner) to improve security. Reconfiguration is used for fault tolerance by dynamically reallocating tasks to different cores in the event of failures [29]. However, the solution focuses on the failure of a core and not on the attack on one task. Besides, the existing solution is not scalable due to its underlying tool (e.g., timed automata).

Researchers explore real-time security problems from various contexts [30]. Hasan et al. propose scheduler-level techniques to integrate security monitoring tasks [25, 26]. However, unlike ours, such approaches do not consider the aftereffect of detecting an anomalous task. There exist architecture-level solutions [11, 31–33] that proactively reset the system to remove malicious entities. As we mentioned in §1, such techniques may not be feasible for some use cases. Further, they need custom hardware and are not designed for periodic hard real-time tasks. In contrast, we propose scheduler-level defense techniques.

There has been research on data and control flow integrity checking for real-time systems [34–38]. However, existing methods operate at the task level, where RESCUE is designed as a scheduler-level defense tool. However, such integrity checkers can be combined with RESCUE to initiate reconfiguration. Some work explores techniques to block potentially vulnerable tasks at a certain point in time, named attack effective window (AEW) [39, 40]. However, they are an offline, proactive defense mechanism. In contrast, RESCUE provides a combination of offline and online solutions and works in a *reactive* manner.

There are also efforts to integrate security mechanisms for both fixed [12, 23, 41] and dynamic priority [42–44] systems. However, they are designed for single-core platforms only. Pre-computed recovery planes for distributed systems can improve the availability of the system in case of a single computing node failure [45]. Neither security nor multi-core aspects are covered in earlier research [45]. To the best of our knowledge, this is one of the first efforts to introduce the notion

of “schedule reconfigurability” to isolate anomalous tasks and thus ensure safety while retaining temporal guarantees in *multi-core* platforms.

7 CONCLUSION

Security threats to critical real-time systems is growing, and traditional security solutions designed for general-purpose systems are often inadequate. Protecting systems that run both safety-critical and non-safety-critical tasks on shared resources, such as multicore platforms, presents a significant challenge. This paper introduces a defense strategy at the scheduler level, named RESCUE, which employs pre-computed recovery plans. Through extensive evaluations, including synthetic workloads and a UAV case study, we demonstrate the efficacy of RESCUE. Our reactive reconfiguration method enhances security and resilience, particularly for multi-core-based time-critical IoT applications and autonomous systems.

ACKNOWLEDGMENT

This research is supported in part by the European Union-funded project CyberSecDome (Agreement 101120779) and the US National Science Foundation (Award 2312006). Any findings, opinions, recommendations, or conclusions expressed in the paper are those of the authors and do not necessarily reflect the views of the sponsors.

REFERENCES

- [1] Z. A. H. Hammadeh, M. Hasan, and M. Hamad, “Securing real-time systems using schedule reconfiguration,” in *2024 IEEE 27th International Symposium on Real-Time Distributed Computing (ISORC)*, 2024, pp. 1–10.
- [2] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, and T. Kohno, “Comprehensive experimental analyses of automotive attack surfaces,” in *20th USENIX security symposium (USENIX Security 11)*, 2011.
- [3] J.-P. Yaacoub, H. Noura, O. Salman, and A. Chehab, “Security analysis of drones systems: Attacks, limitations, and recommendations,” *Internet of Things*, vol. 11, p. 100218, 2020.
- [4] T. Bonaci, J. Yan, J. Herron, T. Kohno, and H. J. Chizeck, “Experimental analysis of denial-of-service attacks on teleoperated robotic systems,” in *Proceedings of the ACM/IEEE sixth international conference on cyber-physical systems*, 2015, pp. 11–20.
- [5] R. Ernst and M. Di Natale, “Mixed criticality systems—a history of misconceptions?” *IEEE Design & Test*, vol. 33, no. 5, pp. 65–74, 2016.
- [6] M. Bechtel and H. Yun, “Memory-aware denial-of-service attacks on shared cache in multicore real-time systems,” *IEEE Transactions on Computers*, vol. 71, no. 9, pp. 2351–2357, 2022.
- [7] C.-Y. Chen, S. Mohan, R. Pellizzoni, R. B. Bobba, and N. Kiyavash, “A novel side-channel in real-time schedulers,” in *2019 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, 2019.
- [8] M. Hamad, M. Tsantekidis, and V. Prevelakis, “Red-zone: Towards an intrusion response framework for intra-vehicle system.” in *VEHITS*, 2019.
- [9] L. Miedema, B. Rouxel, and C. Grelck, “Task-level redundancy vs instruction-level redundancy against single event upsets in real-time dag scheduling,” in *2021 IEEE 14th International Symposium on Embedded Multicore/Many-core Systems-on-Chip (MCSoc)*, 2021, pp. 373–380.
- [10] E. Dubrova, *Fault-tolerant design*. Springer, 2013.
- [11] 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 *2018 ACM/IEEE 9th International Conference on Cyber-Physical Systems (ICCPs)*. IEEE, 2018, pp. 10–21.
- [12] M. Hamad, Z. A. Hammadeh, S. Saidi, V. Prevelakis, and R. Ernst, “Prediction of abnormal temporal behavior in real-time systems,” in *Proceedings of the 33rd Annual ACM Symposium on Applied Computing*, 2018, pp. 359–367.
- [13] 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 *2017 IEEE/ACM Second International Conference on Internet-of-Things Design and Implementation (IoTDI)*, 2017, pp. 191–196.
- [14] C.-Y. Chen, M. Hasan, and S. Mohan, “Securing real-time Internet-of-things,” *Sensors*, vol. 18, no. 12, 2018.
- [15] J. P. Lehoczky, “Fixed priority scheduling of periodic task sets with arbitrary deadlines,” in *IEEE Real-Time Systems Symposium*, 1990.

- [16] G. Bernat, A. Burns, and A. Liamosi, "Weakly hard real-time systems," *IEEE transactions on Computers*, vol. 50, no. 4, pp. 308–321, 2001.
- [17] H. Choi, H. Kim, and Q. Zhu, "Toward practical weakly hard real-time systems: A job-class-level scheduling approach," *IEEE Internet of Things Journal*, vol. 8, no. 8, pp. 6692–6708, 2021.
- [18] E. Bini and G. Buttazzo, "Measuring the performance of schedulability tests," *Real-Time Systems*, vol. 30, no. 1-2, 2005.
- [19] A. Gujarati, F. Cerqueira, and B. B. Brandenburg, "Outstanding paper award: Schedulability analysis of the linux push and pull scheduler with arbitrary processor affinities," in *2013 25th Euromicro Conference on Real-Time Systems*, July 2013, pp. 69–79.
- [20] M. Bertogna and M. Cirinei, "Response-time analysis for globally scheduled symmetric multiprocessor platforms," in *28th IEEE International Real-Time Systems Symposium (RTSS 2007)*, 2007, pp. 149–160.
- [21] Z. A. H. Hammadeh, S. Quinton, and R. Ernst, "Weakly-hard real-time guarantees for earliest deadline first scheduling of independent tasks," *ACM Trans. Embed. Comput. Syst.*, vol. 18, no. 6, dec 2019. [Online]. Available: <https://doi.org/10.1145/3356865>
- [22] M. Hamad, A. Finkenzerler, M. Kühr, A. Roberts, O. Maennel, V. Prevelakis, and S. Steinhorst, "REACT: Autonomous intrusion response system for intelligent vehicles," *Comput. Secur.*, vol. 145, no. C, Nov. 2024. [Online]. Available: <https://doi.org/10.1016/j.cose.2024.104008>
- [23] M. Hasan, S. Mohan, R. B. Bobba, and R. Pellizzoni, "Exploring opportunistic execution for integrating security into legacy hard real-time systems," in *2016 IEEE Real-Time Systems Symposium (RTSS)*. IEEE, 2016, pp. 123–134.
- [24] M. Hasan, S. Mohan, R. Pellizzoni, and R. B. Bobba, "Contego: An Adaptive Framework for Integrating Security Tasks in Real-Time Systems," *LIPICs, Volume 76, ECRTS 2017*, vol. 76, pp. 23:1–23:22, 2017, artwork Size: 22 pages, 885035 bytes ISBN: 9783959770378 Medium: application/pdf Publisher: Schloss Dagstuhl – Leibniz-Zentrum für Informatik. [Online]. Available: <https://drops.dagstuhl.de/entities/document/10.4230/LIPICs.ECRTS.2017.23>
- [25] M. Hasan, S. Mohan, R. Pellizzoni, and R. B. Bobba, "A design-space exploration for allocating security tasks in multicore real-time systems," in *2018 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2018, pp. 225–230.
- [26] M. Hasan, S. Mohan, R. Pellizzoni, and R. B. Bobba, "Period adaptation for continuous security monitoring in multicore real-time systems," in *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2020, pp. 430–435.
- [27] M. Hasan, S. Mohan, R. B. Bobba, and R. Pellizzoni, "Beyond Just Safety: Delay-aware Security Monitoring for Real-time Control Systems," *ACM Transactions on Cyber-Physical Systems*, vol. 6, no. 3, pp. 1–25, Jul. 2022. [Online]. Available: <https://dl.acm.org/doi/10.1145/3520136>
- [28] M. Hasan and S. Mohan, "You can't always check what you wanted: Selective checking and trusted execution to prevent false actuations in real-time internet-of-things," in *2023 IEEE 26th International Symposium on Real-Time Distributed Computing (ISORC)*. IEEE, 2023, pp. 42–53.
- [29] M. T. B. Waez, A. Wąsowski, J. Dingel, and K. Rudie, "Synthesis of a reconfiguration service for mixed-criticality multi-core systems: An experience report," in *Formal Aspects of Component Software: 11th International Symposium, FACS 2014, Bertinoro, Italy, September 10-12, 2014, Revised Selected Papers 11*. Springer, 2015, pp. 162–180.
- [30] M. Hasan, A. Kashinath, C.-Y. Chen, and S. Mohan, "SoK: Security in real-time systems," *ACM Comput. Surv.*, Feb. 2024.
- [31] F. Abdi, C.-Y. Chen, M. Hasan, S. Liu, S. Mohan, and M. Caccamo, "Preserving physical safety under cyber attacks," *IEEE Internet of Things Journal*, vol. 6, no. 4, pp. 6285–6300, 2018.
- [32] L. Niu, D. Sahabandu, A. Clark, and R. Poovendran, "Verifying safety for resilient cyber-physical systems via reactive software restart," in *2022 ACM/IEEE 13th International Conference on Cyber-Physical Systems (ICCPs)*. IEEE, 2022, pp. 104–115.
- [33] R. Romagnoli, B. H. Krogh, D. de Niz, A. D. Hristozov, and B. Sinopoli, "Runtime system support for CPS software rejuvenation," *IEEE Transactions on Emerging Topics in Computing*, 2023.
- [34] N. Bellec, G. Hiet, S. Rokicki, F. Tronel, and I. Puaut, "RT-DFI: Optimizing data-flow integrity for real-time systems," in *ECRTS 2022-34th Euromicro Conference on Real-Time Systems*, no. 34, 2022, pp. 1–24.
- [35] Y. Wang, A. Li, J. Wang, S. Baruah, and N. Zhang, "Opportunistic data flow integrity for real-time cyber-physical systems using worst case execution time reservation," in *Proceedings of the 33rd USENIX Conference on Security Symposium*. USENIX Association, 2024.
- [36] R. J. Walls, N. F. Brown, T. Le Baron, C. A. Shue, H. Okhravi, and B. C. Ward, "Control-flow integrity for real-time embedded systems," in *31st Euromicro Conference on Real-Time Systems (ECRTS 2019)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2019.
- [37] Y. Du, Z. Shen, K. Dharsee, J. Zhou, R. J. Walls, and J. Criswell, "Holistic control-flow protection on real-time embedded systems with kage," in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 2281–2298.
- [38] F. A. T. Abad, J. Van Der Woude, Y. Lu, S. Bak, M. Caccamo, L. Sha, R. Mancuso, and S. Mohan, "On-chip control flow integrity check for real time embedded systems," in *2013 IEEE 1st International Conference on Cyber-Physical Systems*,

- Networks, and Applications (CPSNA)*. IEEE, 2013, pp. 26–31.
- [39] J. Chen, T. Kloda, R. Tabish, A. Bansal, C.-Y. Chen, B. Liu, S. Mohan, M. Caccamo, and L. Sha, “Schedguard++: Protecting against schedule leaks using Linux containers on multi-core processors,” *ACM Transactions on Cyber-Physical Systems*, vol. 7, no. 1, pp. 1–25, 2023.
- [40] J. Ren, C. Liu, C. Lin, R. Bi, S. Li, Z. Wang, Y. Qian, Z. Zhao, and G. Tan, “Protection window based security-aware scheduling against schedule-based attacks,” *ACM Transactions on Embedded Computing Systems*, vol. 22, no. 5s, pp. 1–22, 2023.
- [41] S. Mohan, M.-K. Yoon, R. Pellizzoni, and R. B. Bobba, “Integrating security constraints into fixed priority real-time schedulers,” *Real-Time Systems*, vol. 52, pp. 644–674, 2016.
- [42] V. Lesi, I. Jovanov, and M. Pajic, “Security-aware scheduling of embedded control tasks,” *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 16, no. 5s, pp. 1–21, 2017.
- [43] S. Baruah, “Security-cognizant real-time scheduling,” in *2022 IEEE 25th International Symposium On Real-Time Distributed Computing (ISORC)*. IEEE, 2022, pp. 1–9.
- [44] F. Raadia, N. Fisher, T. Chantem, and S. Baruah, “An improved security-cognizant scheduling model,” in *2024 IEEE 27th International Symposium on Real-Time Distributed Computing (ISORC)*. IEEE, 2024, pp. 1–8.
- [45] A. Lund, Z. A. Haj Hammadeh, P. Kenny, V. Vishav, A. Kovalov, H. Watolla, A. Gerndt, and D. Lüdtkke, “ScOSA system software: the reliable and scalable middleware for a heterogeneous and distributed on-board computer architecture,” *CEAS Space Journal*, vol. 14, no. 1, pp. 161–171, 2022.