

Multimode Security-Aware Real-Time Scheduling on Multiprocessors

Jiankang Ren^{1b}, Member, IEEE, Chunxiao Liu^{1b}, Chi Lin^{1b}, Senior Member, IEEE,
Wei Jiang^{1b}, Senior Member, IEEE, Pengfei Wang^{1b}, Member, IEEE, Xiangwei Qi^{1b},
Simeng Li^{1b}, and Shengyu Li^{1b}

I. INTRODUCTION

Abstract—Embedded real-time systems generally execute in a predictable and deterministic manner to deliver critical functionality within stringent timing constraints. However, the predictable execution behavior leaves the system vulnerable to schedule-based attacks. In this article, we present a multimode security-aware real-time scheduling scheme to counteract schedule-based attacks on multiprocessor real-time systems. To mitigate the vulnerability to the schedule-based attack, we propose a multimode scheduling method to reduce the accumulative attack effective window (AEW) of multiple victim tasks and prevent the untrusted tasks from executing during the AEW by distinctively scheduling mixed-trust tasks according to the system mode. To avoid the protection degradation due to the excessive blocking of untrusted tasks, we introduce a protection window for multiple victims on multiprocessors by analyzing the system protection capability limit under the system schedulability constraint. Furthermore, to maximize the protection capability of the multimode security-aware scheduling strategy on a multiprocessor platform, we also propose a security-aware packing algorithm to balance the workloads of mixed-trust tasks on different processors using a mixed-trust worst-fit decreasing heuristic strategy. The experimental results demonstrate that our proposed approach significantly outperforms the state-of-the-art method. Specifically, the AEW ratio and the AEW untrusted execution time ratio are reduced by 18.8% and 62.8%, respectively, while the defense success rate against ScheduLeak attack is improved by 16.3%.

Index Terms—Multimode scheduling, multiprocessor, real-time systems, schedule-based attacks, security-aware scheduling.

Manuscript received 5 August 2024; accepted 10 August 2024. This work was supported in part by the National Natural Science Foundation of China under Grant 62072067, Grant 62172069, Grant 62072076, Grant 62466059, and Grant 61602080; in part by the Shandong Provincial Natural Science Foundation under Grant ZR2023LZH016; and in part by the Xinjiang Network Information Science and Technology Innovation Research Project under Grant 12421604. This article was presented at the International Conference on Compilers, Architectures, and Synthesis for Embedded Systems (CASES) 2024 and appeared as part of the ESWEEK-TCAD Special Issue. This article was recommended by Associate Editor S. Dailey. (Corresponding author: Chi Lin.)

Jiankang Ren, Chunxiao Liu, Pengfei Wang, Simeng Li, and Shengyu Li are with the Key Laboratory of Social Computing and Cognitive Intelligence, Ministry of Education, Dalian University of Technology, Dalian 116024, China (e-mail: rjk@dlut.edu.cn; liuchuanxiao@mail.dlut.edu.cn; wangpf@dlut.edu.cn; lsmhf@mail.dlut.edu.cn; lisanling@mail.dlut.edu.cn).

Chi Lin is with the School of Software Technology, Dalian University of Technology, Dalian 116024, China (e-mail: c.lin@dlut.edu.cn).

Wei Jiang is with the School of Information and Software Engineering, University of Electronic Science and Technology of China, Chengdu 610054, China (e-mail: weiji Jiang@uestc.edu.cn).

Xiangwei Qi is with the School of Computer Science and Technology, Xinjiang Normal University, Xinjiang 830054, China (e-mail: xiangweiqi10@163.com).

Digital Object Identifier 10.1109/TCAD.2024.3445260

EMBEDDED real-time systems are extensively applied in safety-critical applications, such as automotive, aviation, and industrial robotics. These systems generally execute in a predictable and deterministic manner to deliver critical functionality within stringent timing constraints. However, such a predictable execution pattern exposes a vulnerability that can be exploited by schedule-based attacks [1]. For example, adversaries can leverage the deterministic execution patterns to infer sensitive scheduling information through timing side-channels [2]. With knowledge of system internals gleaned from such attacks, malicious attackers can craft more effective tailored attacks, such as compromising system stability through the injection of inaccurate data [3] and affecting vehicle operations by obstructing control system signals [2], [4]. Moreover, the relentless demand for computational power has driven embedded real-time systems toward multiprocessor architectures. Consequently, it is extremely important to offer an effective security strategy to defend against schedule-based attacks for multiprocessor real-time systems without compromising real-time performance constraints.

The effectiveness of schedule-based attacks typically depends on whether the attacker task is executed during the attack effective windows (AEWs) of victim tasks [5]. For example, the experimental evidence has shown that the AEW for a control output overwrite attack on a customized rover system with real-time Linux is 8.3 ms [2]. According to the timing relationships between the execution states of the victim and the attacker, schedule-based attacks can be divided into four categories [6]: 1) the posterior attack that is launched after the completion of the victim task; 2) the anterior attack that is carried out before the victim task is executed; 3) the concurrent attack that takes place while the victim task is being executed; and 4) the pincer attack that is a hybrid attack combining both posterior and anterior attacks. In this article, we focus on mitigating the posterior schedule-based attack. Note that, as a common threat to real-time systems, the posterior schedule-based attacks, such as replay attacks, zero dynamics attacks, and bias injection attacks, can be launched to manipulate the output of a task, and their attack effects have been demonstrated in Real-time Linux and FreeRTOS [2], [3].

There is a wealth of literature on security-aware scheduling methods to defend against schedule-based attacks (outlined in Section II). Traditionally, research in this field has

73 focused mainly on uniprocessor systems [5], [7], [8], [9],
 74 [10], but lately, attention has turned toward multiprocessor
 75 systems [11], [12]. More recently, Chen et al. [11] proposed a
 76 temporal isolation-based protection approach *SchedGuard++*
 77 to protect against posterior schedule-based attacks on mul-
 78 tiprocessors with multiple victims by preventing untrusted
 79 tasks from running within AEWs of victim tasks. It can
 80 provide best-effort protection for the multiprocessor system
 81 under schedulability constraints. However, it has four major
 82 drawbacks.

- 83 1) It tends to allocate all untrusted tasks to the same
 84 processors, thereby reducing opportunities to miti-
 85 gate schedule-based attacks by strategically scheduling
 86 untrusted tasks with parallel processing capabilities of
 87 multiprocessors.
- 88 2) Upon the completion of a victim task's execution on a
 89 processor, it blocks all other processors, including those
 90 running trusted tasks. This resource wastage leads to a
 91 decline in the performance of security protections under
 92 the system schedulability constraint.
- 93 3) It employs an empirical protection window to guide
 94 the security-aware scheduling, neglecting the system
 95 protection capability limit imposed by schedulability
 96 constraints. This may result in over blocking of untrusted
 97 tasks, ultimately compromising the performance of the
 98 security protection.
- 99 4) It conducts an offline analysis for the maximum tolerable
 100 blocking times of tasks, ignoring the run-time system
 101 behavior. The inherent pessimism of offline analysis
 102 results in a degradation of protection performance,
 103 particularly when handling task sets with high
 104 utilization.

105 To overcome the above limitations, we propose a multimode
 106 security-aware real-time scheduling approach called MM-
 107 SARTS for multiprocessor real-time systems to optimize
 108 the system protection capability under system schedulability
 109 constraints. MM-SARTS enables the system to operate in
 110 different modes, each with its own specific security-aware
 111 scheduling strategy, to mitigate schedule-based attacks by
 112 distinctively scheduling mixed-trust tasks according to the
 113 system's operational status. The primary contributions of our
 114 work can be summarized as follows.

- 115 1) We presented an example to illustrate the limitations of
 116 the isolation-based protection method *SchedGuard++*
 117 for multiprocessor systems with multiple victims.
- 118 2) We proposed a multimode security-aware real-time
 119 scheduling method to mitigate schedule-based attacks
 120 by distinctively scheduling mixed-trust tasks according
 121 to the system mode, coupled with an online priority
 122 inversion feasibility test.
- 123 3) We introduced a protection window for multiple victims
 124 on multiprocessors to avoid protection degradation due
 125 to excessive blocking of untrusted tasks by analyzing
 126 the system protection capability limit under the system
 127 schedulability constraint.
- 128 4) We developed a security-aware task-to-processor pack-
 129 ing algorithm that maximizes the protection capability
 130 of the multimode security-aware scheduling strategy on

a multiprocessor system by balancing the workloads of
 mixed-trust tasks across different processors.
 The experimental evaluation based on an automotive bench-
 mark indicates that our method can effectively decrease the
 AEW ratio and the AEW untrusted execution time ratio and
 enhance the attack defense success rate.

II. RELATED WORK

Since the pioneer research by Son et al. [13] on
 information leakage in real-time systems caused by pre-
 dictable system execution patterns, numerous studies have
 focused on security-aware real-time scheduling strategies to
 counter schedule-based attacks. These security-aware sched-
 uling techniques can generally be categorized into two groups:
 1) randomization-based scheduling and 2) isolation-based
 scheduling.

In randomization-based scheduling, obfuscation mecha-
 nisms are used to diversify the schedule, making it difficult
 for attackers to accurately predict the timing behavior of victim
 tasks [7], [8], [9], [14], [15], [16], [17]. For the fixed-priority
 real-time system in which each task is assigned a static priority
 level, Yoon et al. [7] introduced a randomized security-aware
 scheduling method based on priority inversion. This method
 leverages statically calculated priority inversion budgets to
 resist the schedule-based side-channel attacks while ensuring
 the real-time performance. To increase randomness in the task
 execution pattern, Yoon et al. [14] utilized runtime information
 at each scheduling decision point to enhance the uncertainty
 of task schedules while ensuring system schedulability. For
 the time-triggered real-time system, where tasks are executed
 according to a predetermined and fixed schedule constructed
 based on the schedulability constraints, Krüger et al. [15], [16]
 analyzed vulnerabilities related to timing inference and mali-
 cious behavior and proposed an online job randomization
 method and an offline schedule-diversification method to
 mitigate timing inference-based attacks. For the dynamic-
 priority real-time system, where task priorities are calculated
 during system execution, Chen et al. [8] introduced a ran-
 domized scheduling method to obscure the earliest deadline
 first scheduling policy with limited priority inversions at run-
 time. This method effectively introduces unpredictability into
 execution patterns of tasks, particularly when the system oper-
 ates under low and medium loads. However, its conservative
 predetermination of task priority inversion limits, without
 considering dynamic run-time system behavior, can lead to
 performance degradation under heavy utilization. To address
 this problem, Ren et al. [9] proposed an enhanced randomized
 scheduling strategy that leverages runtime system information
 to increase feasible priority inversion opportunities while
 ensuring system schedulability. Although randomization-based
 scheduling methods can significantly increase timing uncer-
 tainty, making it harder to predict task execution states, it has
 been demonstrated that they may fail to protect against certain
 schedule-based attacks and, in some cases, even increase the
 attack success rate [6].

In isolation-based scheduling, various temporal isolation
 mechanisms are employed to prevent interference among tasks,

187 thereby protecting sensitive information from unauthorized
 188 access [5], [10], [18], [19], [20]. For the fixed-priority real-
 189 time systems, Völöp et al. [18] suggested using an idle
 190 system thread to defend against information leakage. Similarly,
 191 Pellizzoni et al. [20] and Mohan et al. [19] suggested the
 192 introduction of flush tasks to prevent the schedule-based
 193 information leakage between low- and high-security tasks.
 194 However, this mechanism introduces significant overhead,
 195 yielding in a poor response time for real-time tasks and
 196 reducing system schedulability. To prevent the execution of
 197 untrusted tasks during the AEWs, Chen et al. [5] proposed
 198 a coverage-oriented scheduling policy to provide determinis-
 199 tic isolation against posterior schedule-based attacks without
 200 affecting schedulability. However, this approach overlooks
 201 the limit of system protection capacity due to schedulability
 202 constraints in security-aware scheduling decisions, potentially
 203 leading to poor security performance from excessive blocking
 204 of untrusted tasks. Moreover, it conducts an offline analysis of
 205 the maximum acceptable blocking time budget, which results
 206 in diminished protection performance when handling task sets
 207 with high utilization due to the pessimism of the offline
 208 analysis. To avoid the excessive blocking of untrusted tasks
 209 and the pessimism of the offline analysis, Ren et al. [10]
 210 proposed a security-aware real-time scheduling scheme based
 211 on the protection window and the online feasibility test.
 212 However, this method is limited to single-core systems with
 213 a single victim task. For multiprocessor platforms with
 214 multiple victims, Chen et al. [11] introduced an approach
 215 named *SchedGuard++*, which extends the coverage-oriented
 216 scheduling approach in [5] with the worst-case response
 217 time analysis for mixed-trust tasks on the multiprocessor.
 218 Nevertheless, it fails to fully exploit the parallelism of mul-
 219 tiprocessor systems to optimize security performance under
 220 schedulability constraints by decreasing the AEW ratio and
 221 the AEW untrusted execution time ratio.

222 In this article, we propose a novel isolation-based security-
 223 aware scheduling approach for multiprocessor systems with
 224 multiple victims. Our approach differs from these existing
 225 methods in that it can effectively reduce the AEW ratio and
 226 the AEW untrusted execution time ratio through multimode
 227 scheduling based on an online priority inversion feasibility
 228 test. Furthermore, our approach enhances the overall system
 229 protection capability of multiprocessor systems by balancing
 230 mixed-trust task workloads across different processors.

231 III. SYSTEM AND ADVERSARY MODEL

232 A. System Model

233 We consider a real-time system comprising N independent
 234 periodic real-time tasks, denoted by $\Gamma = \{\tau_1, \dots, \tau_N\}$, exe-
 235 cuted on a multiprocessor system with P identical processors
 236 of unit capacity, identified as $\Pi = \{\pi_1, \dots, \pi_P\}$, following
 237 a partitioned fixed-priority preemptive scheduling strategy
 238 commonly used in OSEK/VDX operating system [21] and
 239 AUTOSAR [22]. Each task $\tau_i \in \Gamma$ is defined as $\tau_i =$
 240 (C_i, T_i, D_i) , where

- 241 1) C_i is the worst-case execution time (WCET) of τ_i ;
- 242 2) T_i is the period of τ_i ;

- 3) D_i is the relative deadline of τ_i .

243 For a task τ_i , we use *utilization* symbolized by U_i to indicate
 244 the ratio C_i/T_i , and use $U_\Gamma = \sum_{\tau_i \in \Gamma} U_i$ to represent the
 245 *total utilization* of the task set Γ . We consider all tasks to
 246 be implicit-deadline tasks, where the deadline D_i is equal to
 247 the period T_i , and they are initially released at the same time
 248 $t = 0$. The hyperperiod of a task set Γ is indicated by H_Γ ,
 249 which is the least common multiple of all task periods for task
 250 set Γ . The job of task τ_i is represented as $\mathcal{J}_{i,j}$, and its release
 251 time is indicated as $r_{i,j}$, which is a member of the infinite set
 252 $\{0, T_i, 2T_i, \dots\}$. For a job $\mathcal{J}_{i,j}$ of task τ_i released at $r_{i,j}$, its
 253 completion time is represented as $f_i(r_{i,j})$. If each task $\tau_i \in \Gamma$
 254 can meet its deadline in the worst-case scenario, the system
 255 is considered schedulable. For each task $\tau_i \in \Gamma$, it has a
 256 unique priority, and its priority is allocated based on the rate
 257 monotonic (RM) policy [23]. The set of tasks with priorities
 258 higher than task τ_i is denoted as $hp(\tau_i)$, while the set of tasks
 259 with lower priorities is indicated as $lp(\tau_i)$. Following [5], the
 260 idle times are treated as instances of an extra idle task in
 261 the system, and the idle task has the lowest priority, infinite
 262 execution time, and infinite period and deadline.
 263

264 B. Adversary Model

265 In this article, we follow the vendor-oriented security model
 266 in [20], where information leakage from a vendor's sensitive
 267 tasks to other vendors' tasks is undesirable. For a system,
 268 high-critical tasks (e.g., engine and brake control tasks in
 269 a self-driving system) are regarded as victim tasks. Given
 270 a victim task set, tasks are classified as trusted (from the
 271 same vendor as a victim task, or an idle task) or untrusted
 272 (all other tasks, which may be attackers). It is assumed that
 273 only untrusted tasks pose an attack risk, and the scheduler is
 274 trustworthy.

275 We consider an attack scenario where an adversary carries
 276 out a posterior schedule-based attack on the victim task by
 277 exploiting external connections on the target platform [6]. It
 278 is assumed that the adversary has taken control of certain
 279 tasks, turning them into attackers within the system, and is
 280 able to modify their control flow. The attacker is unaware
 281 of the concrete scheduling scheme, but it can deduce certain
 282 scheduling parameters by monitoring the execution windows
 283 of compromised tasks. For example, such an attack can
 284 covertly deduce the locations and routes of self-driving cars
 285 through the external network [4]. We assume that for the attack
 286 to be effective, it must be carried out within a certain time
 287 window after the victim task completes to steal, corrupt, or
 288 overwrite the victim's output. We define such a time window
 289 as the AEW.

290 *Definition 1 [11]:* For a victim task τ_i^V , its AEW, denoted
 291 by ω_i , is defined as the time period during which schedule-
 292 based attacks are effective and ineffective otherwise.

293 To characterize the amount of AEWs generated by a victim
 294 task within a time interval, we define the accumulative AEWs
 295 of a victim task as follows.

296 *Definition 2 [11]:* Given a scheduling policy \mathcal{P} and a
 297 schedulable task set Γ under \mathcal{P} , for a victim task $\tau_i^V \in \Gamma$, its
 298 *accumulative AEW* within the time interval $[t_1, t_2)$, denoted

by $\Omega(\{\tau_i^v\}, \mathcal{P}, t_1, t_2)$, is defined as the set of all time intervals that belong to the AEW of task τ_i^v within time interval $[t_1, t_2)$.

We focus on multiprocessor real-time systems with multiple victim tasks, and the AEWs of different victim tasks can potentially overlap because of the parallel execution on different processors. Here, we formally define the accumulative AEW for a set of victim tasks as the union of their AEWs.

Definition 3 [11]: Given a scheduling policy \mathcal{P} and a schedulable task set Γ under \mathcal{P} , for the victim task set $\Gamma^v \subseteq \Gamma$, its *accumulative AEW* within the time interval $[t_1, t_2)$, denoted by $\Omega(\Gamma^v, \mathcal{P}, t_1, t_2)$, is defined as the union of AEWs of all victim tasks in Γ^v over the time interval $[t_1, t_2)$, i.e.,

$$\Omega(\Gamma^v, \mathcal{P}, t_1, t_2) = \bigcup_{\tau_i^v \in \Gamma^v} \Omega(\{\tau_i^v\}, \mathcal{P}, t_1, t_2) \quad (1)$$

where $\Omega(\{\tau_i^v\}, \mathcal{P}, t_1, t_2)$ is the accumulative AEW of task τ_i^v over the time interval $[t_1, t_2)$ under scheduling policy \mathcal{P} .

To assess the protection performance of a scheduling policy, we define the AEW ratio and the AEW untrusted execution time ratio as follows.

Definition 4: Given a scheduling policy \mathcal{P} and a schedulable task set Γ under \mathcal{P} , the *AEW ratio* of the task set Γ under \mathcal{P} within the time interval $[t_1, t_2)$, denoted by $\Lambda(\Gamma, \mathcal{P}, t_1, t_2)$, is defined as

$$\Lambda(\Gamma, \mathcal{P}, t_1, t_2) = \frac{L(\Gamma^v, \mathcal{P}, t_1, t_2)}{t_2 - t_1} \quad (2)$$

where Γ^v represents all victim tasks in Γ , and $L(\Gamma^v, \mathcal{P}, t_1, t_2)$ is the cumulative length of the accumulative AEW $\Omega(\Gamma^v, \mathcal{P}, t_1, t_2)$ of task set Γ^v within the time interval $[t_1, t_2)$.

From Definition 4, we can observe that for a task set Γ , a higher AEW ratio indicates a larger attack surface, thereby increasing its susceptibility to attacks.

Definition 5: Given a scheduling policy \mathcal{P} and a schedulable task set Γ under \mathcal{P} , the *AEW untrusted execution time ratio* for the task set Γ within the time interval $[t_1, t_2)$, denoted by $\Theta(\Gamma, \mathcal{P}, t_1, t_2)$, is defined as the total execution time of untrusted tasks within AEWs in $[t_1, t_2)$ divided by their cumulative execution time within the time interval $[t_1, t_2)$, i.e.,

$$\Theta(\Gamma, \mathcal{P}, t_1, t_2) = \frac{\sum_{\tau_i \in \Gamma^u} E_i^{\text{AEW}}(\mathcal{P}, t_1, t_2)}{\sum_{\tau_i \in \Gamma^u} E_i(\mathcal{P}, t_1, t_2)} \quad (3)$$

where Γ^u is the set of all untrusted tasks in Γ , $E_i^{\text{AEW}}(\mathcal{P}, t_1, t_2)$ represents the total execution time of the untrusted task $\tau_i \in \Gamma^u$ within AEWs in $[t_1, t_2)$ under scheduling policy \mathcal{P} and $E_i(\mathcal{P}, t_1, t_2)$ denotes the cumulative execution time of task τ_i within $[t_1, t_2)$ under scheduling policy \mathcal{P} .

From Definition 5, we can see that for a task set Γ , as the AEW untrusted execution time ratio increases, the execution time of untrusted tasks within AEWs tends to be longer, thus resulting in a higher likelihood of successful attacks.

Goal: The main objective of this work is to develop an efficient security-aware multiprocessor real-time scheduling strategy to schedule a given set of periodic real-time tasks on the multiprocessor platform with multiple victim tasks, such that all tasks are schedulable and the chance for the adversary to launch a successful posterior scheduler-based

TABLE I
TASK PARAMETERS OF AN EXAMPLE TASK SET Γ

Task	C_i	T_i	D_i	Trust Level	ω_i
τ_1^v	1	10	10	Trusted	3
τ_2^v	2	20	20	Trusted	5
τ_1^t	1	10	10	Trusted	-
τ_2^t	2	10	10	Trusted	-
τ_3^t	4	20	20	Trusted	-
τ_4^t	4	20	20	Trusted	-
τ_1^u	4	10	10	Untrusted	-
τ_2^u	4	10	10	Untrusted	-
τ_3^u	5	20	20	Untrusted	-
τ_4^u	5	20	20	Untrusted	-

attack is reduced by minimizing the AEW ratio and the AEW untrusted execution time ratio.

IV. MULTIMODE SECURITY-AWARE REAL-TIME SCHEDULING

A. Motivation

Before presenting our multimode security-aware real-time scheduling scheme, we first provide an example to discuss the limitations of the existing temporal isolation-based protection method *SchedGuard++* [11] and to motivate the multimode security-aware real-time scheduling strategy. The basic idea of *SchedGuard++* can be succinctly described as follows.

- 1) In the task-to-processor allocation process, each victim task is initially allocated to a single processor. Next, trusted tasks are evenly distributed among processors with victim tasks based on their utilizations. Finally, untrusted tasks are distributed evenly among processors without victim tasks. It is important to note that, to assign a feasible processor to each task, trusted and untrusted tasks may be allocated to any processor, regardless of whether that processor has a victim task.
- 2) It schedules tasks on the basis of an empirical protection window. Once a victim task is completed on a processor, the protection window starts. During the protection window, all other processors are attempted to be blocked. Note that when a task reaches its maximum tolerable blocking time (i.e., the longest duration that a task can be paused or delayed by lower priority tasks under the schedulability constraint), it is allowed to execute within the protection window to ensure system schedulability.

It has been demonstrated that *SchedGuard++* can defend against posterior schedule-based attacks on multiprocessors by preventing untrusted tasks from being executed during the AEW. However, *SchedGuard++* cannot effectively reduce the AEW ratio and the AEW untrusted execution time ratio for some sets of tasks. Now, we illustrate this with an example.

Example 1: Consider a task set Γ depicted in Table I scheduled on four processors $\Pi = \{\pi_1, \pi_2, \pi_3, \pi_4\}$ under *SchedGuard++*. As shown in Fig. 1, it is the simulation of a synchronous arrival sequence (SAS) for the task set Γ under *SchedGuard++*. From Fig. 1, we can see that all untrusted tasks are allocated to processors π_3 and π_4 , although there remains some capacity on processors π_1 and π_2 . This reduces opportunities to mitigate the schedule-based attack




Fig. 1. SAS simulation for task set Γ under *SchedGuard++*.

393 by strategically scheduling untrusted tasks with the parallel
 394 processing characteristics of multiprocessors. We also can
 395 observe that once a job of victim task τ_1^v is completed, all
 396 other processors, including the processor π_2 with trusted tasks,
 397 are attempted to be blocked. This results in a reduction in
 398 the system's parallel processing capability and an increase in
 399 the AEW ratio. Obviously, some untrusted jobs are executed
 400 during AEWs of victim tasks τ_1^v and τ_2^v . From Fig. 1, we can
 401 see that the accumulative AEW within $[0, 20)$ is $3 + 5 + 3 =$
 402 11 , and hence the AEW ratio within $[0, 20)$ is $11/20 = 0.55$.
 403 We also can see that the total execution time of untrusted
 404 tasks within AEWs in $[0, 20)$ is $2 + 6 = 8$ and the cumulative
 405 execution time of untrusted tasks within $[0, 20)$ is $2 \times 4 + 2 \times$
 406 $4 + 5 + 5 = 26$, and thus the AEW untrusted execution time
 407 ratio within $[0, 20)$ is $8/26 \approx 0.3077$. Note that it is possible
 408 to decrease the AEW ratio within $[0, 20)$ to 0.45 and the AEW
 409 untrusted execution time ratio within $[0, 20)$ to 0 using our
 410 multimode security-aware real-time scheduling approach (see
 411 Example 2).

412 B. Multimode System Model

413 To mitigate the vulnerability to schedule-based attacks,
 414 we model the real-time system as a multimode system and
 415 introduce an enforcement mechanism to prevent untrusted
 416 tasks from being executed during AEWs of victim tasks by
 417 scheduling specific jobs to execute based on the system mode.

418 *Multimode System:* The real-time system is modeled as a
 419 multimode system characterized by a set of system modes \mathcal{M} ,
 420 an initial system mode $M_0 \in \mathcal{M}$, a set of mode transitions
 421 $\mathcal{R} \subseteq \mathcal{M} \times \mathcal{M}$, and a set of implicit-deadline periodic tasks
 422 \mathcal{T} that execute in the system modes. Here, we consider three
 423 system modes: 1) normal mode M^N ; 2) victim mode M^V ; and
 424 3) protection mode M^P , with the initial mode being the normal
 425 mode (i.e., $M_0 = M^N$). For each mode $M \in \mathcal{M}$, we consider
 426 that all tasks in the set \mathcal{T} should be executed. The mode
 427 transition is triggered by a mode change request event (MCR).

428 *Enforcement in a Mode:* To mitigate schedule-based attacks
 429 by strategically scheduling mixed-trust tasks based on the
 430 system's operational mode, we define the scheduling enforce-
 431 ment for each system mode as follows.

432 1) *Normal Mode M^N :* Untrusted tasks are executed as
 433 much as possible on each processor. This enables the
 434 system to execute untrusted tasks to the greatest extent
 435 possible before the execution of victim tasks, preventing
 436 untrusted tasks from launching posterior schedule-based
 437 attacks.

2) *Victim Mode M^V :* Victim tasks are executed as much
 438 as possible on each processor. This allows victim tasks
 439 on different processors to execute simultaneously, max-
 440 imizing the overlap of AEWs of victim tasks across
 441 processors and thereby reducing the accumulative AEW
 442 size of multiple victim tasks in the multiprocessor
 443 system.
 444 3) *Protection Mode M^P :* Trusted and idle tasks are executed
 445 as much as possible on each processor. With this mode,
 446 we can minimize the risk of potential attacks on victim
 447 tasks by preventing untrusted tasks from being executed
 448 during AEWs of victim tasks.
 449

Note that the mode enforcement is *nonstrict*. When a task
 450 must be executed for schedulability, regardless of its type,
 451 it can be executed in any mode to maintain the system
 452 schedulability.
 453

Enforcement Upon an MCR: When the system is operating
 454 in mode M_s and receives an MCR associated with an outgoing
 455 transition (M_s, M_t) , it immediately switches to the new mode
 456 M_t and performs the enforcement. We consider four mode
 457 transitions.
 458

- 459 1) $M^N \rightarrow M^V$: This transition allows the system to execute
 460 untrusted tasks prior to the execution of victim tasks. It
 461 is triggered when no untrusted tasks are pending on any
 462 processor, or when there is a victim task that must be
 463 executed for schedulability.
- 464 2) $M^V \rightarrow M^P$: This transition enables the system to execute
 465 trusted tasks after the completion of victim tasks. It
 466 is triggered when no victim tasks are pending on any
 467 processor.
- 468 3) $M^P \rightarrow M^N$: This transition is used to avoid the
 469 protection degradation due to the excessive blocking of
 470 untrusted tasks. It is triggered when the duration of
 471 the protection mode exceeds a specific threshold. This
 472 threshold is set based on the system protection capability
 473 limit characterized by the protection window given in
 474 Section IV-C.
- 475 4) $M^P \rightarrow M^V$: This transition is designed to ensure the
 476 schedulability of the victim tasks and provide timely
 477 protection for them. It is triggered when a victim task
 478 must be executed for schedulability.

479 C. Protection Window of Multiple Victims on Multiprocessor

480 To effectively mitigate schedule-based attacks by preventing
 481 the protection degradation caused by the excessive blocking
 482 of untrusted tasks in the protection mode, we introduce the
 483 protection window for multiple victim tasks on the mul-
 484 tiprocessor platform. This protection window characterizes
 485 the limit of system protection capability under schedulability
 486 constraints.

487 *Definition 6:* Given a schedulable task set $\Gamma(\pi_k)$ on a
 488 uniprocessor π_k , the protection window of a victim job $\mathcal{J}_{i,j}^v$ of
 489 victim task $\tau_i^v \in \Gamma(\pi_k)$, denoted by $S_{i,j}$, is defined as the length
 490 of time interval $[f^v(r_{i,j}^v), \min\{r_{i,j}^v + T_i^v, b^u\})$, where $f^v(r_{i,j}^v)$ is
 491 the finish time of the job $\mathcal{J}_{i,j}^v$, T_i^v is the period of task τ_i^v ,
 492 and b^u is the start time of execution of the first untrusted job
 493 executed after $f^v(r_{i,j}^v)$.

494 *Definition 7:* Given a schedulable task set $\Gamma(\pi_k)$ on a
 495 uniprocessor π_k , the protection window of a victim task $\tau_i^v \in$
 496 $\Gamma(\pi_k)$, denoted by \mathcal{S}_i , is defined as the minimum protection
 497 window of jobs for victim task τ_i^v , i.e.,

$$498 \quad \mathcal{S}_i = \min_{1 \leq j \leq H_{\Gamma(\pi_k)}/T_i^v} \{S_{i,j}\} \quad (4)$$

499 where $S_{i,j}$ is the protection window of job $\mathcal{J}_{i,j}^v$ for τ_i^v , $H_{\Gamma(\pi_k)}$
 500 is the hyperperiod of $\Gamma(\pi_k)$, and T_i^v is the period of τ_i^v .

501 *Lemma 1:* Given a schedulable task set $\Gamma(\pi_k)$ on a unipro-
 502 cessor π_k , the protection window of a victim task $\tau_i^v \in \Gamma(\pi_k)$
 503 is bounded by

$$504 \quad \mathcal{S}_i^{\max} = T_i^v \left(1 - \left(\sum_{\tau_j \in \Gamma^u(\pi_k)} U_j \right) - U_i^v \right) \quad (5)$$

505 where $\Gamma^u(\pi_k)$ is the set of all untrusted tasks in task set $\Gamma(\pi_k)$,
 506 T_i^v is the period of task τ_i^v , and U_j and U_i^v represent the
 507 utilizations of tasks τ_j and τ_i^v , respectively.

508 *Proof:* For a hyperperiod with $m = H_{\Gamma(\pi_k)}/T_i^v$ jobs of the
 509 victim task τ_i^v , where $H_{\Gamma(\pi_k)}$ is the hyperperiod of task set
 510 $\Gamma(\pi_k)$ and T_i^v is the period of task τ_i^v , we can derive the
 511 following inequality based on Definition 7:

$$512 \quad m\mathcal{S}_i \leq \sum_{1 \leq l \leq m} S_{i,l} \quad (6)$$

513 where $S_{i,l}$ is the protection window of victim job $\mathcal{J}_{i,l}^v$, and \mathcal{S}_i
 514 is the protection window of victim task τ_i^v .

515 According to Definition 6, there is no execution of any
 516 untrusted task or victim task τ_i^v within the protection window
 517 of the job of victim task τ_i^v , and hence we can derive

$$518 \quad \sum_{1 \leq l \leq m} S_{i,l} \leq H_{\Gamma(\pi_k)} - \left(\sum_{\tau_j \in \Gamma^u(\pi_k)} \frac{H_{\Gamma(\pi_k)}}{T_j} C_j \right) - \frac{H_{\Gamma(\pi_k)}}{T_i^v} C_i^v$$

$$519 \quad = mT_i^v \left(1 - \left(\sum_{\tau_j \in \Gamma^u(\pi_k)} U_j \right) - U_i^v \right). \quad (7)$$

520 From (6) and (7), we can derive the following:

$$521 \quad \mathcal{S}_i \leq T_i^v \left(1 - \left(\sum_{\tau_j \in \Gamma^u(\pi_k)} U_j \right) - U_i^v \right). \quad (8)$$

522 Thus, the protection window of the victim task $\tau_i^v \in \Gamma(\pi_k)$
 523 is bounded by \mathcal{S}_i^{\max} given in (5). ■

524 *Definition 8:* Given a schedulable task set $\Gamma(\pi_k)$ on a
 525 uniprocessor π_k , the protection window of all victim tasks in
 526 task set $\Gamma(\pi_k)$, denoted by $\mathcal{S}(\pi_k)$, is defined as the minimum
 527 protection window of victim tasks in task set $\Gamma(\pi_k)$, i.e.,

$$528 \quad \mathcal{S}(\pi_k) = \min_{\tau_i^v \in \Gamma^v(\pi_k)} \{\mathcal{S}_i\} \quad (9)$$

529 where $\Gamma^v(\pi_k)$ is the set of all victim tasks in task set $\Gamma(\pi_k)$,
 530 and \mathcal{S}_i is the protection window of victim task $\tau_i^v \in \Gamma^v(\pi_k)$.

531 Note that for a processor without a victim task, we assume
 532 its protection window to be infinite. Based on Definition 8 and
 533 Lemma 1, we can directly derive the following lemma.

Algorithm 1 Multimode Security-Aware Scheduling

Input: Scheduling point t , current system mode M_t , task set $\Gamma(\pi_k)$,
 ready queue Q_t , victim ready queue Q_t^v , trusted ready queue
 Q_t^T , untrusted ready queue Q_t^U .

Output: A job executed at time t .

```

1: if  $M_t = M^N$  then
2:    $\mathcal{J}_{\text{test}} \leftarrow$  the highest priority job in  $Q_t^U$ 
3: else if  $M_t = M^V$  then
4:    $\mathcal{J}_{\text{test}} \leftarrow$  the highest priority job in  $Q_t^v$ 
5: else
6:   if  $Q_t^T \neq \emptyset$  then
7:      $\mathcal{J}_{\text{test}} \leftarrow$  the highest priority job in  $Q_t^T$ 
8:   else
9:      $\mathcal{J}_{\text{test}} \leftarrow$  idle job
10:  end if
11: end if
12: if  $\mathcal{J}_{\text{test}}$  is the highest priority job in  $Q_t$  then
13:    $\mathcal{J}_{\text{test}}$  is executed until the next scheduling point.
14: else
15:    $(\text{flag}, E_{\text{test}}^t) \leftarrow$  FeasibilityTest( $\Gamma(\pi_k), \mathcal{J}_{\text{test}}, t$ )
16:   if  $\text{flag} = \text{True}$  then
17:      $\mathcal{J}_{\text{test}}$  is executed until the next scheduling point.
18:   else
19:     The highest priority job in  $Q_t$  is executed until the next
       scheduling point.
20:   end if
21: end if

```

534 *Lemma 2:* Given a schedulable task set $\Gamma(\pi_k)$ on a unipro-
 535 cessor π_k , the protection window $\mathcal{S}(\pi_k)$ of all victim tasks in
 536 task set $\Gamma(\pi_k)$ is bounded by

$$537 \quad \mathcal{S}^{\max}(\pi_k) = \min_{\tau_i^v \in \Gamma^v(\pi_k)} \{\mathcal{S}_i^{\max}\} \quad (10)$$

538 where $\Gamma^v(\pi_k)$ is the set of all victim tasks in task set $\Gamma(\pi_k)$,
 539 and \mathcal{S}_i^{\max} is the protection window upper bound of task $\tau_i^v \in$
 540 $\Gamma^v(\pi_k)$ given in (5).

541 *Definition 9:* Given a schedulable task set Γ scheduled on
 542 a multiprocessor Π with a partitioned scheduling policy, the
 543 protection window of all victim tasks in the task set Γ , denoted
 544 by \mathcal{S}_{Π} , is defined as follows:

$$545 \quad \mathcal{S}_{\Pi} = \min_{\pi_k \in \Pi} \{\mathcal{S}(\pi_k)\} \quad (11)$$

546 where $\mathcal{S}(\pi_k)$ is the protection window of all victim tasks in
 547 the task set $\Gamma(\pi_k)$ assigned to processor $\pi_k \in \Pi$.

548 Based on Definition 9 and Lemma 2, the following theorem
 549 can be directly derived.

550 *Theorem 1:* Given a schedulable task set Γ scheduled on
 551 a multiprocessor Π with a partitioned scheduling policy, the
 552 protection window of all victim tasks in Γ is bounded by

$$553 \quad \mathcal{S}_{\Pi}^{\max} = \min_{\pi_k \in \Pi} \{\mathcal{S}^{\max}(\pi_k)\} \quad (12)$$

554 where $\mathcal{S}^{\max}(\pi_k)$ is the protection window upper bound of tasks
 555 assigned to the processor $\pi_k \in \Pi$ given in (10).

556 For a task set Γ scheduled on a multiprocessor Π with a
 557 partitioned scheduling policy, once the duration of the pro-
 558 tection mode M^P exceeds \mathcal{S}_{Π}^{\max} , the system will immediately
 559 enter the normal mode M^N to avoid the protection degradation
 560 due to excessive blocking of untrusted tasks.

561 D. Multimode Security-Aware Scheduling

562 Based on the multimode system model given in Section IV-B
 563 we propose a multimode security-aware real-time scheduling
 564 algorithm to provide best-effort security protection for victim
 565 tasks while maintaining system schedulability. The key idea of
 566 this scheduling algorithm is to reduce the accumulative AEW of
 567 multiple victim tasks and prevent untrusted tasks from executing
 568 during the AEW by distinctively scheduling mixed-trust tasks
 569 according to the system mode.

570 As given in Algorithm 1, it is our multimode security-aware
 571 scheduling algorithm written in pseudocode. In the algorithm,
 572 there is a ready queue Q_t that holds all jobs that are ready
 573 to run and waiting to be executed, with Q_t^V , Q_t^T , and Q_t^U
 574 representing the ready queues for victim jobs, trusted jobs,
 575 and untrusted jobs, respectively. The scheduling decisions are
 576 made by selecting a candidate job according to the system
 577 mode. When the system is in normal mode, the execution
 578 feasibility of the ready untrusted job with the highest priority
 579 will be checked (lines 1 and 2). When the system is in victim
 580 mode, the execution feasibility of the ready victim job with
 581 the highest priority will be checked (lines 3 and 4). When the
 582 system mode is in protection mode, if there are some trusted
 583 jobs in the ready queue, the execution feasibility of the ready
 584 trusted job with the highest priority will be checked (lines 6
 585 and 7); otherwise, the execution feasibility of the idle job will
 586 be checked (line 9). Note that it is always feasible for the
 587 highest priority job in the ready queue Q_t , since there is no
 588 priority inversion for its execution (lines 12 and 13). If $\mathcal{J}_{\text{test}}$
 589 is not the highest priority job in ready queue Q_t , an online
 590 priority inversion feasibility test is performed for $\mathcal{J}_{\text{test}}$, and the
 591 online feasibility test algorithm is given in Algorithm 2. If it
 592 is feasible to execute the selected job, the selected job will be
 593 executed at t (lines 16 and 17); otherwise, the highest priority
 594 job in the ready queue Q_t will be executed (line 19). Note that
 595 the selected job is executed until the next scheduling point.
 596 In our scheduling algorithm, the scheduling points include the
 597 arrival time of the MCR associated with a mode translation,
 598 the completion time of the current job, the moment when the
 599 running job experienced the feasible execution time E_{test}^t and
 600 the release instant of new jobs.

601 From Algorithm 1, we can find that some higher priority
 602 tasks can be blocked by lower priority tasks (i.e., priority
 603 inversion) during the security-aware scheduling of each system
 604 mode. To ensure system schedulability in the presence of pri-
 605 ority inversion, we conduct an online feasibility test based on
 606 the priority inversion budget analysis, similar to the approach
 607 in [10]. For convenience, we provide a summary of the online
 608 feasibility test based on priority inversion budget analysis here.

609 **Definition 10 [7]:** Given a task set $\Gamma(\pi_k)$ scheduled on a
 610 uniprocessor π_k , the maximum priority inversion budget of
 611 task $\tau_h \in \Gamma(\pi_k)$ at time t , denoted by \mathcal{B}_h^t , is defined as the
 612 maximum amount of time for which all lower priority tasks
 613 $lp(\tau_h)$ are allowed to execute before τ_h finishes at t , while
 614 ensuring the schedulability of τ_h in the worst-case scenario.

615 In the analysis of the priority inversion budget for a task
 616 $\tau_h \in hp(\tau_{\text{test}})$, we consider its two adjacent jobs \mathcal{J}_h^t and $\mathcal{J}_h^{t'}$,
 617 where job \mathcal{J}_h^t is the last job of task τ_h released no later than
 618 time t and job $\mathcal{J}_h^{t'}$ is the first job of task τ_h released after

Algorithm 2 Online Feasibility Test

Input: Task set $\Gamma(\pi_k)$, \bar{V}_i of $\tau_i \in \Gamma(\pi_k)$ calculated off-line with
 (19) and (20), test job $\mathcal{J}_{\text{test}}$, scheduling point t , ready queue Q_t ,
 job \mathcal{J}_h^t that is the last job of task $\tau_h \in hp(\tau_{\text{test}})$ released no later
 than time t .
Output: The feasibility *flag* of executing job $\mathcal{J}_{\text{test}}$ at time t , and
 the maximum feasible execution time E_{test}^t of job $\mathcal{J}_{\text{test}}$.

```

1: flag  $\leftarrow$  True
2:  $E_{\text{test}}^t \leftarrow \tilde{C}_{\text{test}}^t$ 
3: for all  $\tau_h \in hp(\tau_{\text{test}})$  do
4:   Calculate  $I_h(t, d_h^t)$  with equation (14)
5:   if  $\mathcal{J}_h^t \in Q_t$  at time  $t$  then
6:      $\mathcal{B}_h^t \leftarrow d_h^t - t - \tilde{C}_h^t - I_h(t, d_h^t)$ 
7:   else
8:      $\mathcal{B}_h^t \leftarrow d_h^t - t + \bar{V}_h - I_h(t, d_h^t)$ 
9:   end if
10:  if  $\mathcal{B}_h^t > 0$  then
11:     $E_{\text{test}}^t \leftarrow \min\{E_{\text{test}}^t, \overline{\mathcal{B}_h^t}\}$ 
12:  else
13:    flag  $\leftarrow$  False
14:     $E_{\text{test}}^t \leftarrow 0$ 
15:    break
16:  end if
17: end for
18: return (flag,  $E_{\text{test}}^t$ )

```

time t . Depending on whether \mathcal{J}_h^t is in the ready queue Q_t or
 not at time t , we consider two cases. 619 620

Case 1 (Job \mathcal{J}_h^t Is in Ready Queue Q_t at Time t): In this
 case, we can focus solely on the analysis of the job \mathcal{J}_h^t of task
 τ_h at time t . Let d_h^t be the absolute deadline of job \mathcal{J}_h^t . Based
 on Definition 10, by analyzing the workload during the time
 interval $[t, d_h^t)$, the maximum priority inversion budget \mathcal{B}_h^t of
 task τ_h at time t can be expressed as 621 622 623 624 625 626

$$\mathcal{B}_h^t \geq d_h^t - t - \tilde{C}_h^t - I_h(t, d_h^t) \quad (13) \quad 627$$

where \tilde{C}_h^t is the worst-case remaining execution time of job
 \mathcal{J}_h^t . $I_h(t, d_h^t)$ is the worst-case workload of tasks in $hp(\tau_h)$
 during the time interval $[t, d_h^t)$, and it can be calculated by 628 629 630

$$I_h(t, d_h^t) = \sum_{\tau_j \in hp(\tau_h)} \tilde{C}_j^t + \sum_{\tau_j \in hp(\tau_h)} \left\lceil \frac{d_h^t - r_j^{t'}}{T_j} \right\rceil C_j \quad (14) \quad 631$$

where $\lceil x \rceil_0$ is the smallest non-negative integer greater than
 or equal to x , \tilde{C}_j^t is the worst-case remaining execution time
 of task τ_j at time t , and $r_j^{t'}$ is the release time of the first job
 of task τ_j released after time t . 632 633 634 635

Case 2 (Job \mathcal{J}_h^t Is Not in Ready Queue Q_t at Time t): In
 this scenario, the upcoming job $\mathcal{J}_h^{t'}$ of task τ_h released after
 time t should be analyzed. Let $r_h^{t'}$ and $d_h^{t'}$ be the release time
 and the absolute deadline of job $\mathcal{J}_h^{t'}$. Based on Definition 10,
 the maximum priority inversion budget \mathcal{B}_h^t of task τ_h at time
 t can be expressed as 636 637 638 639 640 641

$$\mathcal{B}_h^t = \mathcal{B}_h(t, r_h^{t'}) + \mathcal{B}_h(r_h^{t'}, d_h^{t'}) \quad (15) \quad 642$$

where $\mathcal{B}_h(t, r_h^{t'})$ and $\mathcal{B}_h(r_h^{t'}, d_h^{t'})$ are the maximum priority
 inversion budgets of task τ_h during time intervals $[t, r_h^{t'})$ and 643 644

645 $[r_h^t, d_h^t)$, respectively. By analyzing the workload during the
646 time interval $[t, r_h^t)$, $\mathcal{B}_h^t(t, r_h^t)$ can be expressed as

$$647 \quad \mathcal{B}_h(t, r_h^t) \geq r_h^t - t - I_h(t, r_h^t) \quad (16)$$

648 where $I_h(t, r_h^t)$ is the worst-case workload of tasks in $hp(\tau_h)$
649 during the time interval $[t, r_h^t)$. Since task τ_h is an implicit-
650 deadline periodic task (i.e., $r_h^t = d_h^t$), (16) can be rewritten as

$$651 \quad \mathcal{B}_h(t, r_h^t) \geq d_h^t - t - I_h(t, d_h^t) \quad (17)$$

652 where $I_h(t, d_h^t)$ can be calculated with (14).

653 For the maximum priority inversion budget $\mathcal{B}_h(r_h^t, d_h^t)$ of
654 task τ_h during time interval $[r_h^t, d_h^t)$, it can be expressed as

$$655 \quad \mathcal{B}_h(r_h^t, d_h^t) \geq \bar{V}_h \quad (18)$$

656 where \bar{V}_h is the maximum amount of time that τ_h can
657 additionally have while meeting its deadline when there are
658 no deferred executions of tasks in $hp(\tau_h)$ at the release time
659 of task τ_h . According to [14], \bar{V}_h can be calculated offline,
660 and it can be expressed as

$$661 \quad \bar{V}_h = \max \{ \delta \mid W_h(\delta) \leq D_h \} \quad (19)$$

662 where D_h is the relative deadline of task τ_h , and $W_h(\delta)$ is the
663 duration between a critical instant and the response completion
664 of the corresponding request of task τ_h with extra execution
665 time δ . According to [23], $W_h(\delta)$ can be derived by solving
666 the following iterative formula:

$$667 \quad W_h^{n+1}(\delta) = W_h^0(\delta) + \sum_{\tau_j \in hp(\tau_h)} \left\lceil \frac{W_h^n(\delta)}{T_j} \right\rceil C_j \quad (20)$$

668 where $W_h^0(\delta)$ is the initial value of $W_h(\delta)$, and it is set as
669 $C_h + \delta$ by considering the WCET and the extra execution time
670 of task τ_h . $W_h^n(\delta)$ is the value of $W_h(\delta)$ at the n th iteration.

671 By (15), (17), and (18), the maximum priority inversion
672 budget \mathcal{B}_h^t of task τ_h at time t can be expressed as

$$673 \quad \mathcal{B}_h^t \geq d_h^t - t - I_h(t, d_h^t) + \bar{V}_h. \quad (21)$$

674 Based on the analysis of the two cases mentioned above,
675 the maximum priority inversion budget \mathcal{B}_h^t of task τ_h at time
676 t is bounded by (13) and (21). Hence, a lower bound of \mathcal{B}_h^t ,
677 denoted by $\bar{\mathcal{B}}_h^t$, can be calculated by

$$678 \quad \bar{\mathcal{B}}_h^t = \begin{cases} d_h^t - t - \tilde{C}_h^t - I_h(t, d_h^t) & \mathcal{J}_h^t \in Q_t \text{ at time } t \\ d_h^t - t + \bar{V}_h - I_h(t, d_h^t) & \mathcal{J}_h^t \notin Q_t \text{ at time } t. \end{cases} \quad (22)$$

679 From (22), for a task τ_{test} , if $\bar{\mathcal{B}}_h^t$ is greater than zero for each
680 task $\tau_h \in hp(\tau_{\text{test}})$, it is feasible to execute task τ_{test} at time
681 t by the priority inversion. Thus, we can obtain the following
682 theorem.

683 **Theorem 2 [10]:** For a schedulable task set $\Gamma(\pi_k)$ on a
684 uniprocessor π_k under the RM policy, if the job $\mathcal{J}_{\text{test}}$ of
685 task $\tau_{\text{test}} \in \Gamma(\pi_k)$ is executed with execution time E_{test}^t at
686 time t when $E_{\text{test}}^t \leq \bar{\mathcal{B}}_h^t$ for all tasks $\tau_h \in hp(\tau_{\text{test}})$ and
687 $\bar{\mathcal{B}}_h^t$ is calculated with (22), then the task set $\Gamma(\pi_k)$ is also
688 schedulable.

689 *Proof:* For the detailed proof, please refer to [10]. ■

As shown in Algorithm 2, it is the priority inversion budget-
690 based online feasibility test algorithm written in pseudocode.
691 In this algorithm, *flag* indicates whether $\mathcal{J}_{\text{test}}$ can be executed
692 at time t and E_{test}^t denotes its maximum feasible execution
693 time at time t . First, *flag* and E_{test}^t are initialized to True
694 and $\tilde{C}_{\text{test}}^t$, respectively (lines 1 and 2). Then, the feasibility
695 of executing the job $\mathcal{J}_{\text{test}}$ at time t is checked by calculating
696 the lower bound of the maximum priority inversion budget for
697 each task $\tau_h \in hp(\tau_{\text{test}})$. The upper bound on the workload
698 of each task in $hp(\tau_h)$ during the interval $[t, d_h^t)$ is calculated
699 by (14) (line 4). The lower bound of the maximum priority
700 inversion budget for task τ_h is calculated by (22) (lines 5–9).
701 If $\bar{\mathcal{B}}_h^t \geq 0$, the job $\mathcal{J}_{\text{test}}$ can be executed at t with execution
702 time $\bar{\mathcal{B}}_h^t$ while ensuring the schedulability of task τ_h , and E_{test}^t
703 is updated to $\min\{E_{\text{test}}^t, \bar{\mathcal{B}}_h^t\}$ (lines 10 and 11). If there exists
704 a task $\tau_h \in hp(\tau_{\text{test}})$ whose $\bar{\mathcal{B}}_h^t$ is not greater than 0, it is not
705 feasible to execute $\mathcal{J}_{\text{test}}$ at time t (lines 12–16); otherwise, it
706 is feasible. 707

Lemma 3: Let $\Gamma(\pi_k)$ be the set of tasks scheduled on the
708 uniprocessor π_k . The execution feasibility of a job $\mathcal{J}_{\text{test}}$ of task
709 $\tau_{\text{test}} \in \Gamma(\pi_k)$ at time t can be obtained with the computational
710 complexity $\mathcal{O}(|hp(\tau_{\text{test}})|^2)$ with Algorithm 2, where $hp(\tau_{\text{test}})$
711 is the set of higher priority tasks of τ_{test} in task set $\Gamma(\pi_k)$ and
712 $|hp(\tau_{\text{test}})|$ is the number of tasks in $hp(\tau_{\text{test}})$. 713

Proof: From Algorithm 2, it is evident that the computa-
714 tional complexity of the priority inversion budget-based online
715 feasibility test for the job $\mathcal{J}_{\text{test}}$ depends mainly on the number
716 of tasks with higher priority than τ_{test} (i.e., line 3) and
717 the complexity to calculate $I_h(t, d_h^t)$ with (14) (i.e., line 4).
718 The number of tasks with higher priority than task τ_{test} is
719 $|hp(\tau_{\text{test}})|$. The complexity to calculate $I_h(t, d_h^t)$ with (14) is
720 $\mathcal{O}(|hp(\tau_h)|)$ for each task $\tau_h \in hp(\tau_{\text{test}})$. Since $|hp(\tau_h)| \leq$
721 $|hp(\tau_{\text{test}})|$ for each task $\tau_h \in hp(\tau_{\text{test}})$, we can obtain that the
722 complexity of the feasibility test with Algorithm 2 for the job
723 $\mathcal{J}_{\text{test}}$ is $\mathcal{O}(|hp(\tau_{\text{test}})|^2)$. ■ 724

Theorem 3: For a schedulable task set $\Gamma(\pi_k)$ on a uniproc-
725 essor π_k under the RM policy, it is also schedulable with
726 the multimode security-aware scheduling strategy given in
727 Algorithm 1. 728

Proof: According to lines 12–21 in Algorithm 1, for each
729 system mode, any job of tasks in $\Gamma(\pi_k)$ is executed with
730 a priority inversion or based on the priority assigned with
731 the RM policy. According to Theorem 2, for each system
732 mode, the schedulability of $\Gamma(\pi_k)$ can be ensured when
733 some jobs are executed with priority inversion based on the
734 online feasibility test given in Algorithm 2, since the task
735 set $\Gamma(\pi_k)$ is schedulable under the RM policy. Thus, mode
736 schedulability can be ensured for all system modes. Moreover,
737 when an MCR associated with a mode translation arrives, the
738 system will immediately enter the new mode and perform the
739 enforcement by calling Algorithm 1, and thus mode transition
740 schedulability can also be ensured. Therefore, we can conclude
741 that the task set $\Gamma(\pi_k)$ is also schedulable under the multimode
742 security-aware scheduling strategy. ■ 743

Theorem 4: Let $\Gamma(\pi_k)$ be the task set scheduled on uniproc-
744 essor π_k with the multimode security-aware scheduling
745 policy given in Algorithm 1. The computational complexity of
746

Algorithm 1 is $\mathcal{O}(|\Gamma(\pi_k)|^2)$, where $|\Gamma(\pi_k)|$ is the number of tasks in the task set $\Gamma(\pi_k)$.

Proof: From Algorithm 1, it is evident that the computational complexity of the multimode security-aware scheduling algorithm primarily relies on the computational complexity of the online feasibility test based on the priority inversion budget analysis given in Algorithm 2, and the feasibility test is conducted no more than once for each scheduling point (lines 12–21). By referring to Lemma 3, it can be deduced that the computational complexity of Algorithm 1 is $\mathcal{O}(|\Gamma(\pi_k)|^2)$, which is polynomial complexity. ■

5.9 E. Security-Aware Task Partitioning

In this section, we present a partitioning algorithm to assign a set of mixed-trust real-time tasks Γ to P identical, unit-capacity processors Π . The objective of this algorithm is to balance the mixed-trust task workloads across processors while ensuring system schedulability, such that the system protection capability can be maximized under the schedulability constraint. This is achieved with a mixed-trust WF decreasing heuristic strategy, which sequentially selects victim tasks, trusted tasks, and untrusted tasks to assign based on the WF decreasing heuristic.

Algorithm 3 is our partitioning algorithm, written in pseudocode. Here, the set of tasks assigned to processor π_k is identified by $\Gamma(\pi_k)$. The algorithm starts by initializing $\Gamma(\pi_k)$ to null (line 1). Then, it tries to assign suitable processors to victim tasks (lines 4–15), trusted tasks (lines 16–29), and untrusted tasks (lines 30–43) in sequence. For each type of tasks, tasks are verified in descending utilization order (lines 4, 16, and 30). In this order, each task is tried on each of the processors in Π , ordered in increasing utilization. For a task, if no feasible processor is found, the algorithm aborts with a failure (lines 13, 27, and 41). If all tasks are successfully assigned, the algorithm reports success (line 44).

Theorem 5: Given an implicit-deadline periodic real-time task set $\Gamma = \{\tau_1, \tau_2, \dots, \tau_N\}$, if it is successfully partitioned by Algorithm 3 on P processors $\Pi = \{\pi_1, \pi_2, \dots, \pi_P\}$ where the feasibility test of each processor is performed based on the schedulability condition of the RM policy, then all tasks in Γ can meet their deadlines under the multimode security-aware scheduling policy given in Algorithm 1.

Proof: We consider any processor $\pi_k \in \Pi$, and the task set $\Gamma(\pi_k)$ assigned to π_k . According to lines 5–10, 19–24, and 33–38 in Algorithm 3, the task set $\Gamma(\pi_k)$ is feasible on processor π_k under the RM scheduling policy. Based on Theorem 3, task set $\Gamma(\pi_k)$ is also feasible under the scheduling strategy given in Algorithm 1. This implies that the task set on each local processor can be successfully scheduled. Consequently, we can conclude that the resulting task allocation obtained by Algorithm 3 always guarantees that all tasks in Γ can meet the deadlines under the multimode security-aware scheduling policy given in Algorithm 1 when task set Γ is successfully partitioned by Algorithm 3. ■

Example 2: Consider the task set Γ given in Table I again. We consider that the task set Γ is allocated to four processors

Algorithm 3 Security-Aware Task Partitioning

Input: Task set $\Gamma = \Gamma^v \cup \Gamma^t \cup \Gamma^u$, multiprocessor platform $\Pi = \{\pi_1, \dots, \pi_P\}$.

Output: Task partitions $\{\Gamma(\pi_1), \Gamma(\pi_2), \dots, \Gamma(\pi_P)\}$.

- 1: $\Gamma(\pi_k) \leftarrow \emptyset$, for all $k = 1, \dots, P$.
- 2: **for** all $\tau_i^v \in \Gamma^v$ in descending order of U_i^v **do**
- 3: $flag \leftarrow \text{False}$
- 4: **for** all $\pi_k \in \Pi$ in increasing order of $U_{\Gamma(\pi_k)}$ **do**
- 5: **if** it is feasible for task set $\Gamma(\pi_k) \cup \{\tau_i^v\}$ **then**
- 6: $\Gamma^v \leftarrow \Gamma^v \setminus \tau_i^v$
- 7: $\Gamma(\pi_k) \leftarrow \Gamma(\pi_k) \cup \{\tau_i^v\}$
- 8: $flag \leftarrow \text{True}$
- 9: **break**
- 10: **end if**
- 11: **end for**
- 12: **if** $flag = \text{False}$ **then**
- 13: **return** FAILURE
- 14: **end if**
- 15: **end for**
- 16: **for** all $\tau_i^t \in \Gamma^t$ in descending order of U_i^t **do**
- 17: $flag \leftarrow \text{False}$
- 18: **for** all $\pi_k \in \Pi$ in increasing order of $U_{\Gamma(\pi_k)}$ **do**
- 19: **if** it is feasible for task set $\Gamma(\pi_k) \cup \{\tau_i^t\}$ **then**
- 20: $\Gamma^t \leftarrow \Gamma^t \setminus \tau_i^t$
- 21: $\Gamma(\pi_k) \leftarrow \Gamma(\pi_k) \cup \{\tau_i^t\}$
- 22: $flag \leftarrow \text{True}$
- 23: **break**
- 24: **end if**
- 25: **end for**
- 26: **if** $flag = \text{False}$ **then**
- 27: **return** FAILURE
- 28: **end if**
- 29: **end for**
- 30: **for** all $\tau_i^u \in \Gamma^u$ in descending order of U_i^u **do**
- 31: $flag \leftarrow \text{False}$
- 32: **for** all $\pi_k \in \Pi$ in increasing order of $U_{\Gamma(\pi_k)}$ **do**
- 33: **if** it is feasible for task set $\Gamma(\pi_k) \cup \{\tau_i^u\}$ **then**
- 34: $\Gamma^u \leftarrow \Gamma^u \setminus \tau_i^u$
- 35: $\Gamma(\pi_k) \leftarrow \Gamma(\pi_k) \cup \{\tau_i^u\}$
- 36: $flag \leftarrow \text{True}$
- 37: **break**
- 38: **end if**
- 39: **end for**
- 40: **if** $flag = \text{False}$ **then**
- 41: **return** FAILURE
- 42: **end if**
- 43: **end for**
- 44: **return** SUCCESS

$\Pi = \{\pi_1, \pi_2, \pi_3, \pi_4\}$ with the security-aware task partitioning algorithm given in Algorithm 3, and the tasks allocated to each processor are scheduled with the multimode security-aware scheduling algorithm given in Algorithm 1. From the SAS simulation for the task set Γ given in Fig. 2, we can see that all tasks are schedulable and that all untrusted jobs are executed outside AEWs of all victim tasks to protect all victim jobs from being attacked by untrusted tasks. We can see that the accumulative AEW within time interval $[0, 20)$ is $6 + 3 = 9$, and hence the AEW ratio within time interval $[0, 20)$ is $9/20 = 0.45$. Since there is no untrusted task execution within AEWs in time interval $[0, 20)$, we can obtain that the AEW untrusted execution time ratio within time interval $[0, 20)$ is 0. Therefore, MM-SARTS can provide better protection than




Fig. 2. SAS simulation for task set Γ under MM-SARTS.

817 *SchedGuard++* by effectively reducing the AEW ratio and
 818 the AEW untrusted execution time ratio with security-aware
 819 task partitioning and multimode security-aware scheduling.

820 V. EVALUATION

821 We conducted experimental evaluations to assess the
 822 performance of our multimode security-aware multiprocessor
 823 real-time scheduling method using synthetically generated
 824 workloads based on automotive benchmark [24]. In the exper-
 825 iment, we implemented the following six algorithms.

- 826 1) *RM-FF*: RM scheduling combining with the first-fit (FF)
 827 bin-packing heuristic algorithm.
- 828 2) *RM-NF*: RM scheduling combining with the next-fit
 829 (NF) bin-packing heuristic algorithm.
- 830 3) *RM-BF*: RM scheduling combining with the best-fit
 831 (BF) bin-packing heuristic algorithm.
- 832 4) *RM-WF*: RM scheduling combining with the WF bin-
 833 packing heuristic algorithm.
- 834 5) *SchedGuard++*: The security-aware multiprocessor
 835 real-time scheduling method from [11].
- 836 6) *MM-SARTS*: Our multimode security-aware multipro-
 837 cessor real-time scheduling method.

838 *Objectives*: Our evaluation has two primary goals: 1)
 839 to compare the attack defense performance of our multi-
 840 mode security-aware multiprocessor scheduling approach with
 841 existing scheduling methods in terms of AEW ratio, AEW
 842 untrusted execution time ratio and ScheduLeak attack defense
 843 effect and 2) to assess the overhead of different scheduling
 844 methods by measuring the scheduler CPU time consump-
 845 tion with `time.process_time_ns()` method in the Python time
 846 module.

847 A. Evaluation Setup

848 The periods of all tasks are automotive specific semi-
 849 harmonic, and they are drawn at random from the set $\{1, 2,$
 850 $5, 10, 20, 50, 100, 200, 1000\}$ with an associated appearance
 851 probability given in [24]. Task utilization is determined with
 852 the UUniFast approach from [25] and task priorities are
 853 assigned according to the RM scheduling policy. We define
 854 the normalized utilization $U_{\text{nor}}(\Gamma)$ of a task set Γ deployed
 855 on the multiprocessor platform $\Pi = \{\pi_1, \dots, \pi_P\}$ to be

$$856 \quad U_{\text{nor}}(\Gamma) = \frac{U_{\Gamma}}{P} \quad (23)$$

857 where U_{Γ} is the total utilization of the task set Γ , and P is the
 858 processor number. To ensure the generated task set's utilization
 859 within a specific narrow range around the desired normalized

utilization, with the above parameters, we generated the tasks 860
 one at a time until the system utilization met the condition 861

$$862 \quad U_{\text{nor}}^* - 0.005 \leq U_{\text{nor}}(\Gamma) \leq U_{\text{nor}}^* + 0.005 \quad (24)$$

where U_{nor}^* ranges from 0.05 to 0.95 with step size of 0.05. 863

864 We consider that there are four processors (i.e., $P = 4$) in
 the system. For each value of U_{nor}^* , we generate 1000 task sets.
 865 For each task set, there are 20–30 tasks. The experimental
 866 results are averaged over these 1000 task sets. 867

868 Following the experimental setting in [11], 40% of the tasks
 869 within the task set are chosen at random to be considered as
 870 trusted tasks. 50% of the trusted tasks are randomly selected
 871 as victim tasks while excluding the lowest priority task in the
 872 task set. The AEWs of the victim tasks are determined as a
 873 percentage from the set $\{10, 30, 50\}$ of the period. 874

875 To validate the performance of MM-SARTS against the
 876 schedule-based attack, we evaluated the defensive capabilities
 877 of different scheduling approaches against the ScheduLeak
 878 attack [2]. ScheduLeak is a common schedule-based attack
 879 that utilizes a system timer to measure and reconstruct the
 880 valid execution intervals of the attacker task with a lower
 881 priority than the victim task. In our experiments, a ScheduLeak
 882 attacker task is chosen at random from the untrusted tasks
 883 in the task set. It is important to note that MM-SARTS can
 884 be applied to the system with multiple attacker tasks. This
 885 is because in MM-SARTS, every untrusted task is viewed
 886 as a potential attacker task in the security-aware scheduling
 887 process, and thus MM-SARTS can offer protection against all
 888 untrusted tasks rather than focusing on a specific untrusted
 889 task. 890

891 The experiments were conducted on a desktop computer
 892 equipped with an AMD Ryzen 7 5800H CPU running at 3.20
 GHz with 8 cores, featuring 16 GB of physical RAM, and
 893 operating on a Linux kernel version 5.15.0-88-generic. 894

895 B. Results

896 *AEW Ratio*: Fig. 3 illustrates the AEW ratio within a
 897 hyperperiod versus the utilization of the task set of differ-
 898 ent algorithms for the systems with different AEW sizes.
 899 Fig. 3 shows that RM-WF has a lower AEW ratio relative
 900 to the other nonsecurity-aware RM scheduling approaches.
 901 The possible reason for this is that RM-WF can effectively
 902 balance the workload across different processors, thus fully
 903 utilizing the system's parallel processing capabilities to reduce
 904 the AEW ratio. We can also observe that *SchedGuard++*
 905 performs worse than all nonsecurity-aware RM scheduling
 906 methods in most cases, especially for the system with short
 907 AEWs. The main reason for this is that once a victim task
 908 finishes, *SchedGuard++* blocks all other processors, including
 909 processors running trusted tasks. This decreases the chance
 that victim tasks execute in parallel, thereby reducing the part
 of AEWs that overlaps between different processors. 910

911 Our main result is that MM-SARTS *consistently outper-*
 912 *forms* all existing scheduling algorithms. The performance
 913 gap tends to widen as the utilization decreases; the reason is
 914 that, as the utilization decreases, there are more opportunities
 for MM-SARTS to reduce the AEW ratio by increasing the 915




Fig. 3. AEW ratio for different algorithms. (a) AEW is 10% of victim task period. (b) AEW is 30% of victim task period. (c) AEW is 50% of victim task period.




Fig. 4. AEW untrusted execution time ratio for different algorithms. (a) AEW is 10% of victim task period. (b) AEW is 30% of victim task period. (c) AEW is 50% of victim task period.

915 AEW overlap with security-aware task-to-processor packing
 916 and multimode security-aware scheduling. Specifically, for
 917 the system with an AEW size of 10%, when the system
 918 utilization is 0.60, MM-SARTS demonstrates an enhancement
 919 of approximately 18.8% over *SchedGuard++*, and a boost of
 920 around 13.2% compared to the best nonsecurity-aware RM
 921 scheduling method RM-WF [see Fig. 3(a)]. From Fig. 3(a), we
 922 can also find that the improvement of MM-SARTS in the AEW
 923 ratio is at its lowest point when the system utilization is 0.95.
 924 Even in this case, MM-SARTS still achieves an improvement
 925 of 11.8% over *SchedGuard++*, and a gain of 7.5% compared
 926 to the RM-WF scheduling method.

927 *AEW Untrusted Execution Time Ratio*: Fig. 4 illustrates the
 928 AEW untrusted execution time ratio within a hyperperiod
 929 versus the system utilization of different algorithms for the
 930 systems with different AEW sizes. It clearly shows that RM-
 931 WF has a lower AEW untrusted execution time ratio compared
 932 to the other three RM scheduling approaches. We can observe
 933 that as the AEW size and system utilization increase, the AEW
 934 untrusted execution time ratios for all algorithms also increase.
 935 The reason is that, as the AEW size and system utilization
 936 increase, the AEW ratio and the untrusted task workload tend
 937 to increase. Note that when the system utilization is 0.05, for
 938 all AEW settings, the AEW untrusted execution time ratios
 939 of all methods tend to approach zero, though there is still a
 940 portion of task sets with a ratio greater than zero. Moreover,
 941 it is notable that *SchedGuard++* outperforms all nonsecurity-
 942 aware RM-WF scheduling methods by effectively preventing
 943 the execution of untrusted tasks during the specified AEW.

944 It is not surprising that MM-SARTS *consistently performs*
 945 *better* than *SchedGuard++* in all scenarios. This is primarily
 946 due to MM-SARTS effectively reducing the AEW untrusted
 947 execution time ratio in a multiprocessor real-time system with
 948 multiple victims by distributing the mixed-trust task workload
 949 evenly across different processors and strategically scheduling
 950 mixed-trust tasks on each processor based on the system
 951 mode. Moreover, the enhancement in the AEW untrusted
 952 execution time ratio of MM-SARTS tends to rise as the system
 953 utilization increases for all AEW sizes. Specifically, in a




Fig. 5. Defense success rate of different algorithms.

954 system with an AEW size of 10%, when the system utilization
 955 is 0.6, MM-SARTS shows a 62.8% improvement compared to
 956 *SchedGuard++* and an 86.8% improvement compared to the
 957 RM-WF scheduling method [see Fig. 4(a)].

958 *ScheduLeak Attack Defense Effect*: Fig. 5 illustrates the
 959 defense success rate against ScheduLeak versus the task set
 960 utilization of different algorithms. Here, the defense success
 961 rate is calculated as the failure rate of ScheduLeak to infer
 962 accurate parameters of the victim task. From Fig. 5, we
 963 can see that RM-WF has a higher defense success rate
 964 relative to the other three RM scheduling approaches. We
 965 also can see that both *SchedGuard++* and MM-SARTS
 966 exhibit superior defense capabilities against attacks compared
 967 to nonsecurity-aware RM-WF scheduling methods. This supe-
 968 riority is attributed to the ability of *SchedGuard++* and
 969 MM-SARTS to prevent the attacker task from running after
 970 the victim task completes, creating a false impression for the
 971 attacker about the victim's execution time. Consequently, the
 972 attacker is misled into launching an attack at an incorrect
 973 time based on inaccurate timing information. Additionally,
 974 it can be observed that MM-SARTS *consistently surpasses*
 975 *SchedGuard++*, as MM-SARTS enhances defense effec-
 976 tiveness through security-aware task-to-processor allocation
 977 and multimode security-aware scheduling based on online
 978 feasibility test. Specifically, when the system utilization is
 979 0.6, MM-SARTS shows a 16.3% improvement compared to
 980 *SchedGuard++*.

981 *Scheduler Overhead*: Fig. 6 illustrates a comparison of the
 982 average online scheduling time within ten hyperperiods versus




Fig. 6. Average online scheduling time of different algorithms.

983 system utilization for various algorithms. It is evident that
 984 the nonsecurity-aware RM scheduling methods have the least
 985 scheduler overhead, as they do not require online feasibility
 986 test for the priority inversion. We also can see that the overhead
 987 of MM-SARTS is slightly higher than that of *SchedGuard++*.
 988 The main reason for this is that MM-SARTS explores more
 989 opportunities for priority inversion to enhance protection
 990 performance, resulting in more online feasibility tests.

991 VI. CONCLUSION

992 We have proposed MM-SARTS, a multimode security-
 993 aware real-time scheduling technique against schedule-based
 994 attacks in fixed-priority real-time systems on multiprocessors.
 995 MM-SARTS works by distinctively scheduling mixed-trust
 996 tasks with an online priority inversion feasibility test according
 997 to system modes to minimize the AEW ratio and the AEW
 998 untrusted execution time ratio. In particular, we introduce the
 999 protection window for multiple victims on multiprocessors
 1000 to avoid the protection degradation due to the excessive
 1001 blocking of untrusted tasks by analyzing the system protection
 1002 capability limit under the system schedulability constraint. To
 1003 maximize the protection capability of the multimode security-
 1004 aware scheduling strategy on the multiprocessor platform,
 1005 we also propose a security-aware task-to-processor packing
 1006 algorithm to balance the workloads of mixed-trust tasks across
 1007 different processors. Our evaluation shows that MM-SARTS
 1008 surpasses the existing security-aware scheduling method for
 1009 fixed-priority real-time systems on multiprocessors in terms of
 1010 the AEW ratio, the AEW untrusted execution time ratio, and
 1011 the attack defense capability.

1012 REFERENCES

1013 [1] D. Trilla et al., “Cache side-channel attacks and time-predictability in
 1014 high-performance critical real-time systems,” in *Proc. 55th Annu. Design
 1015 Autom. Conf.*, 2018, pp. 1–6.
 1016 [2] C.-Y. Chen, S. Mohan, R. Pellizzoni, R. B. Bobba, and N. Kiyavash,
 1017 “A novel side-channel in real-time schedulers,” in *Proc. IEEE Real-
 1018 Time Embed. Technol. Appl. Symp. (RTAS)*, Montreal, QC, Canada, 2019,
 1019 pp. 90–102.
 1020 [3] S. Bi and Y. J. Zhang, “False-data injection attack to control
 1021 real-time price in electricity market,” in *Proc. GLOBECOM*, 2013,
 1022 pp. 772–777.

[4] M. Luo et al., “Stealthy tracking of autonomous vehicles with
 1023 cache side channels,” in *Proc. 29th USENIX Secur. Symp.*, 2020,
 1024 pp. 859–876.
 1025 [5] J. Chen et al., “SchedGuard++: Protecting against schedule leaks using
 1026 Linux containers,” in *Proc. IEEE 27th Real-Time Embed. Technol. Appl.
 1027 Symp.*, 2021, pp. 14–26.
 1028 [6] M. Nasri, T. Chantem, G. Bloom, and R. M. Gerdes, “On the pitfalls
 1029 and vulnerabilities of schedule randomization against schedule-based
 1030 attacks,” in *Proc. IEEE Real-Time Embedd. Technol. Appl. Symp.
 1031 (RTAS)*, 2019, pp. 103–116.
 1032 [7] M.-K. Yoon, S. Mohan, C.-Y. Chen, and L. Sha, “TaskShuffler:
 1033 A schedule randomization protocol for obfuscation against tim-
 1034 ing inference attacks in real-time systems,” in *Proc. RTAS*, 2016,
 1035 pp. 1–12.
 1036 [8] C.-Y. Chen, M. Hasan, A. Ghassami, S. Mohan, and N. Kiyavash,
 1037 “REORDER: Securing dynamic-priority real-time systems using sched-
 1038 ule obfuscation,” 2018, *arXiv:1806.01393*.
 1039 [9] J. Ren et al., “REORDER++: Enhanced randomized real-time schedul-
 1040 ing strategy against side-channel attacks,” *IEEE Trans. Netw. Sci. Eng.*,
 1041 vol. 10, no. 6, pp. 3253–3266, Nov./Dec. 2023.
 1042 [10] J. Ren et al., “Protection window based security-aware scheduling
 1043 against schedule-based attacks,” *ACM Trans. Embed. Comput. Syst.*,
 1044 vol. 22, no. 5s, pp. 1–22, 2023.
 1045 [11] J. Chen et al., “SchedGuard++: Protecting against schedule leaks using
 1046 Linux containers on multi-core processors,” *ACM Trans. Cyber-Phys.
 1047 Syst.*, vol. 7, no. 1, pp. 1–25, 2023.
 1048 [12] M. Hasan, S. Mohan, R. Pellizzoni, and R. B. Bobba, “Period
 1049 adaptation for continuous security monitoring in multicore real-time
 1050 systems,” in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, 2020,
 1051 pp. 430–435.
 1052 [13] J. Son et al., “Covert timing channel analysis of rate monotonic real-
 1053 time scheduling algorithm in MLS systems,” in *Proc. IEEE Inf. Assur.
 1054 Workshop*, 2006, pp. 361–368.
 1055 [14] M.-K. Yoon, J.-E. Kim, R. Bradford, and Z. Shao, “TaskShuffler++:
 1056 Real-time schedule randomization for reducing worst-case vulnerability
 1057 to timing inference attacks,” 2019, *arXiv:1911.07726*.
 1058 [15] K. Krüger, M. Völpl, and G. Fohler, “Vulnerability analysis and mit-
 1059 igation of directed timing inference based attacks on time-triggered
 1060 systems,” in *Proc. 30th ECRTS*, 2018, pp. 1–22.
 1061 [16] K. Krüger et al., “Randomization as mitigation of directed timing
 1062 inference based attacks on time-triggered real-time systems with task
 1063 replication,” *Leibniz Trans. Embed. Syst.*, vol. 7, no. 1, pp. 01:1–01:29,
 1064 2021.
 1065 [17] M.-K. Yoon, J.-E. Kim, R. Bradford, and Z. Shao, “TimeDice:
 1066 Schedulability-preserving priority inversion for mitigating covert timing
 1067 channels between real-time partitions,” in *Proc. 52nd DSN*, 2022,
 1068 pp. 453–465.
 1069 [18] M. Völpl, C.-J. Hamann, and H. Härtig, “Avoiding timing channels in
 1070 fixed-priority schedulers,” in *Proc. ACM Symp. Inf., Comput. Commun.
 1071 Secur.*, 2008, pp. 44–55.
 1072 [19] S. Mohan, M. K. Yoon, R. Pellizzoni, and R. B. Bobba, “Real-time
 1073 systems security through scheduler constraints,” in *Proc. 26th ECRTS*,
 1074 2014, pp. 129–140.
 1075 [20] R. Pellizzoni et al., “A generalized model for preventing information
 1076 leakage in hard real-time systems,” in *Proc. 21st RTAS*, 2015,
 1077 pp. 271–282.
 1078 [21] *Road Vehicles Open Interface for Embedded Automotive Applications
 1079 Part 3: OSEK/VDX Operating System*, ISO Standard 17356-3,
 1080 2005.
 1081 [22] “Specification of operating system,” AUTOSAR Consortium,
 1082 Hörgertshausen, Germany, Rep. CP R22-11, 2022.
 1083 [23] C. L. Liu et al., “Scheduling algorithms for multiprogramming in a
 1084 hard-real-time environment,” *J. ACM*, vol. 20, no. 1, pp. 46–61, 1973.
 1085 [24] S. Kramer, D. Ziegenbein, and A. Hamann, “Real world automotive
 1086 benchmarks for free,” in *Proc. WATERS*, 2015, pp. 1–6.
 1087 [25] E. Bini and G. C. Buttazzo, “Measuring the performance of schedula-
 1088 bility tests,” *Real-Time Syst.*, vol. 30, nos. 1–2, pp. 129–154, 2005.
 1089