# Period Adaptation for Continuous Security Monitoring in Multicore Real-Time Systems

Monowar Hasan<sup>\*</sup>, Sibin Mohan<sup>\*</sup>, Rodolfo Pellizzoni<sup>†</sup>, Rakesh B. Bobba<sup>‡</sup> {mhasan11, sibin}@illinois.edu, <sup>†</sup>rodolfo.pellizzoni@uwaterloo.ca, <sup>‡</sup>rakesh.bobba@oregonstate.edu

# ABSTRACT

We propose a design-time framework (named HYDRA-C) for integrating security tasks into partitioned<sup>1</sup> real-time systems (RTS) running on multicore platforms. Our goal is to opportunistically execute security monitoring mechanisms in a 'continuous' manner - i.e., as often as possible, across cores, to ensure that security tasks run with as few interruptions as possible. Our framework will allow designers to integrate security mechanisms without perturbing existing real-time (RT) task properties or execution order. We demonstrate the framework using a proof-of-concept implementation with intrusion detection mechanisms as security tasks. We develop and use both, (a) a custom intrusion detection system (IDS), as well as (b) Tripwire - an open source data integrity checking tool. These are implemented on a realistic rover platform designed using an ARM multicore chip. We compare the performance of HYDRA-C with a state-of-the-art RT security integration approach for multicore-based RTS and find that our method can, on average, detect intrusions 19.05% faster without impacting the performance of RT tasks.

### **1** INTRODUCTION

Limited resources in terms of processing power, memory, energy, etc. coupled with the fact that security was not considered a design priority has led to the deployment of a large number of realtime systems (RTS) that include little to no security mechanisms. Hence, retrofitting such legacy RTS with general-purpose security solutions is a challenging problem since any perturbation of the real-time (RT) constraints (runtimes, periods, task execution orders, deadlines, etc.) could be detrimental to the correct and safe operation of RTS. Besides, security mechanisms need to be designed in such a way that an adversary can not easily evade them. Successful attacks/intrusions into RTS are often aimed at impacting the safety guarantees of such systems, as evidenced by recent intrusions (e.g., attacks on control systems [2, 3], automobiles [4, 5], medical devices [6], etc. to name but a few). Systems with RT properties pose unique security challenges - these systems are required to meet stringent timing requirements along with strong safety requirements. Limited resources (i.e., computational power, storage, energy, etc.) prevent security mechanisms that have been primarily developed for general purpose systems from being effective for safety-critical RTS.

In this paper we aim to improve the security posture of RTS through integration of security tasks while ensuring that the existing RT tasks are not affected by such integration. The security

#### **Table 1: Example of Security Tasks**

Security Task	Approach/Tools
File-system checking	Tripwire [17], AIDE [18], etc.
Network packet monitoring	Bro [19], Snort [20], etc.
Hardware event monitoring	Statistical analysis based
	checks [21] using performance
	monitors (e.g., perf [22],
	OProfile [23], etc.)
Application specific	Behavior-based detection (see
checking	the related work [11–13, 24])

tasks considered could be carrying out any one of protection, detection or response-based operations, depending on the system requirements. For instance, a sensor measurement correlation task may be added for detecting sensor manipulation or a change detection task (or other intrusion detection programs) may be added to detect changes/intrusions into the system. In Table 1 we present some examples of security tasks that can be integrated into legacy systems (this is by no stretch meant to be an exhaustive list). Note that the addition of any security mechanisms (such as IDS, encryption/authentication, behavior-based monitoring, *etc.*) may require modification of the system or the RT task parameters as was the case in prior work [7–15].

Further, to provide the best protection, *security tasks may need to be executed as often as possible.* If the interval between consecutive checking events is too large then an attacker may remain undetected and cause harm to the system between two invocations of the security task. In contrast, if the security tasks are executed very frequently, it may impact the schedulability of the RT (and other security) tasks. The challenge is then to determining the right *periods* (*i.e.*, minimum inter-invocation time) for the security tasks [16].

As a step towards enabling the design of secure RT platforms, opportunistic execution [25, 26] has been proposed as a potential way to integrate security mechanisms into legacy RTS – this allows the execution of security mechanisms as background services without impacting the timing constraints of the RT tasks. Other approaches have been built on this technique for integrating tasks into both legacy and non-legacy systems [7–11, 27, 28]. However, most of that work was focused on single core RTS (that are a majority of such systems in use today). However, *multicore* processors have found increased use in the design of RTS to improve overall performance and energy efficiency [29, 30]. While the use of such processors *increases* the security problems in RTS (*e.g.*, due to parallel execution of critical tasks) [31] to our knowledge very few security solutions have been proposed in literature [26]. In prior work (called HYDRA) [26] researchers have developed a

<sup>&</sup>lt;sup>1</sup>In *partitioned scheduling* (a widely accepted multicore scheduling scheme), tasks are statically partitioned onto identical cores (*i.e.*, runtime migration across cores is not permitted) [1].

mechanism for integrating security into multicore RTS. However this work uses a partitioned scheduling approach and does not allow runtime migration of security tasks across cores. We show that this results in delayed detection of intrusions<sup>2</sup> as the security tasks are not able to execute as frequently. Our main goal in this paper is to raise the responsiveness of such security tasks by increasing their frequency of execution. For instance, consider an intrusion detection system (IDS) - say one that checks the integrity of file systems. If such a system is interrupted (before it can complete checking the entire system), then an adversary could use that opportunity to intrude into the system and, perhaps, stay resident in the part of the filesystem that has already been checked (assuming that the IDS is carrying out the check in parts). If, on the other hand, the IDS task is able to execute with as few interruptions as possible (e.g., by moving immediately to an empty core when it is interrupted), then there is much higher chance of success and, correspondingly, a much lower chance of a successful adversarial action.

*Our Contributions.* In this paper, we propose a design-time methodology and a framework named HYDRA-C for partitioned<sup>3</sup> RTS that (*a*) leverages semi-partitioned scheduling [35] to enable *continuous execution* of security tasks (*i.e.*, execute as frequently as possible) across cores, and (*b*) does not impact the timing constraints of other, existing, RT tasks.

HYDRA-C takes advantage of the properties of a multicore platform and allows security tasks to migrate across available cores and execute opportunistically (i.e., when the RT tasks are not running). This framework extends existing work [26] and ensures better security (e.g., faster detection time) and schedulability (see Section 5). HYDRA-C is able to do this without violating timing constraints for either the existing RT tasks or the security ones (Section 3). We develop a mathematical model and iterative solution that allows security tasks to execute as frequently as possible while still considering the schedulability constraints of other tasks (Section 4). In addition, we also present an implementation on a realistic ARM-based multicore rover platform (running a RT variant of Linux system and realistic security applications). We then perform comparisons with the state-of-the-art [26] (Section 5.1). Finally, we carry out a design space exploration using synthetic workloads and study trade-offs for schedulability and security. Our evaluation shows that proposed semi-partitioned approach can achieve better execution frequency for security tasks and consequently quicker intrusion detection (19.05% faster on average) when compared with both fully-partitioned and global scheduling approaches while providing the same or better schedulability (Section 5.2).

**Note:** We do not target our framework towards any specific security mechanism – our focus is to integrate any designer-provided security solution into a multicore-based RTS. In our experiments we used Tripwire [17] (a data integrity checking tool) as well as our *in-house custom-developed malicious kernel module checker* to demonstrate the feasibility of our approach –

the integration framework proposed in this paper is more broadly applicable to other security mechanisms.

# 2 MODEL AND ASSUMPTIONS

# 2.1 Real-time Tasks and Scheduling Model

Consider a set of  $N_R$  RT tasks  $\Gamma_R = \{\tau_1, \tau_2, \cdots, \tau_{N_R}\}$ , scheduled on a multicore platform with M identical cores  $\mathcal{M} = \{\pi_1, \pi_2, \cdots, \tau_M\}$ . Each RT task  $\tau_r$  releases an infinite sequence of task instances, also called *jobs*, and is represented by the tuple  $(C_r, T_r, D_r)$  where  $C_r$  is the worst-case execution time (WCET),  $T_r$  is the minimum inter-arrival time (*e.g.*, period) and  $D_r$  is the relative deadline. The utilization of each task is denoted by  $U_r = \frac{C_r}{T_r}$ . We assume constrained deadlines for RT tasks (*e.g.*,  $D_r \leq T_r$ ) and that the task priorities are assigned according to rate-monotonic (RM) [36] order (*e.g.*, shorter period implies higher priority).

All events in the system happen with the precision of integer clock ticks (*i.e.*, processor clock cycles), that is, any time *t* involved in scheduling is a non-negative integer. In this paper we consider RT tasks that are scheduled using partitioned fixed-priority preemptive scheme [30] and assigned to the cores using a standard task partitioning algorithm [1, 30]. We further assume that the RT tasks are *schedulable*, *viz.*, the worst-case response time (WCRT), denoted as  $\mathcal{R}_r$ , is less than deadline (*e.g.*,  $\mathcal{R}_r \leq D_r$ ,  $\forall \tau_r$ ) and the following necessary and sufficient schedulability condition holds for each RT tasks  $\tau_r$  assigned to any given core  $\pi_m$  [1]:

$$\exists t : 0 < t \le D_r \text{ and } C_r + \sum_{\tau_i \in hp(\tau_r, \pi_m)} \left\lceil \frac{t}{T_i} \right\rceil C_i \le t, \tag{1}$$

where  $hp(\tau_r, \pi_m)$  denotes the set of RT tasks with higher priority than  $\tau_r$  assigned to core  $\pi_m$ .

# 2.2 Security Model

Our focus is on integrating given security mechanisms abstracted as security tasks into a legacy multicore RTS without impacting the RT functionality of the RTS. While we use specific intrusion detection mechanisms (*e.g.*, Tripwire) to demonstrate our approach, our approach is somewhat agnostic to the security mechanisms. The security model used and the design of security tasks are orthogonal problems. Since we aim to maximize the frequency of execution of security tasks, security mechanisms whose performance improves with frequency of execution (*e.g.*, intrusion monitoring and detection tasks) benefit from our framework.

# **3 SECURITY INTEGRATION FRAMEWORK**

We propose to improve the security posture of multicore based RT systems by integrating additional *periodic security tasks* (*e.g.*, tasks that are specifically designed for intrusion detection purposes). We highlight that HYDRA-C abstracts security tasks and allows designers to execute *any given techniques*. Our focus here is on integration of a given set of security tasks (*e.g.*, intrusion detection mechanisms) in an existing multicore RTS without impacting the RT task parameters (*e.g.*, WCET, periods, *etc.*) or their task execution order. In general, the addition of security mechanisms may increase the execution time of existing tasks [7, 8] or reduce schedulability [15]. As we mentioned earlier, our focus is on legacy

<sup>&</sup>lt;sup>2</sup>We discuss this issue further in Section 5.

<sup>&</sup>lt;sup>3</sup>Since this is the commonly used multicore scheduling approach for many commercial and open-source OSs (such as OKL4 [32], QNX [33], RT-Linux [34], *etc.*) – mainly due to its simplicity and efficiency [1, 26].



Figure 1: Illustration of our security integration framework for a dual-core platform: two RT tasks (blue and green) are statically assigned to two cores (core 0 and core 1, respectively). We propose to integrate a security task (red) that will execute with lowest priority and can be migrated to ether core (whichever is idle) at runtime.

multicore systems where designers may not have enough flexibility to modify system parameters to integrate security mechanisms. We address this problem by allowing security tasks to execute with a priority lower than all the RT tasks, i.e., leverage opportunistic execution [25, 26]. This way, security tasks will only execute during the slack time (e.g., when a core is idle) and the timing requirements of the RT tasks will not be perturbed. However, in contrast to prior work (HYDRA) [26] where the security tasks are statically bound to their respective cores, in this paper we allow security tasks to continuously migrate at runtime (i.e., the combined taskset with RT and security tasks follows a semi-partitioned scheduling policy) whenever any core is available (e.g., when other RT or higherpriority security tasks are not running). An illustration of HYDRA-C is presented in Fig. 1 where two RT tasks (represented by blue and green rectangles) are partitioned into two cores and a newly added security task (red rectangle) can move across cores.

As we shall see in Section 5, allowing security tasks to execute on any available core will give us the opportunity to execute security tasks more frequently (*e.g.*, with shorter period) and that leads to better responsiveness (faster intrusion detection time). One fundamental question with our security integration approach is to figure out how often to execute security tasks so that the system remains schedulable (*e.g.*, WCRT is less than period), and also can execute within a designer provided frequency bound (so that the security checking remains effective). This is different when compared to scheduling traditional RT tasks since the RT task parameters (*e.g.*, periods) are often derived from physical system properties and cannot be adjusted due to control/application requirements. We now formally define security tasks.

Security Tasks. Let us include a set of  $N_S$  security tasks  $\Gamma_S = \{\tau_1, \tau_2, \cdots, \tau_{N_S}\}$  in the system. We adopt the periodic security task model [25] and represent each security task by the tuple  $(C_s, T_s, T_s^{max})$  where  $C_s$  is the WCET,  $T_s$  is the (unknown) period (e.g.,  $\frac{1}{T_s}$  is the monitoring frequency) and  $T_s^{max}$  is a designer provided upper bound of the period – if the period of the security task is higher than  $T_s^{max}$  then the responsiveness is too low and security checking may not be effective.

We assume that priority of the security tasks are distinct and specified by the designers (*e.g.*, derived from specific security requirements). Security tasks have implicit deadlines, *i.e.*, they need to finish execution before the next invocation. We also assume that task migration and context switch overhead is negligible compared to WCET of the task. Our goal here is to find a minimum period  $T_s \leq T_s^{max}$  (so that the security tasks can execute more frequently) such that the taskset remains schedulable (*e.g.*,  $\forall \tau_s \in \Gamma_S: \mathcal{R}_s \leq T_s$  where  $\mathcal{R}_s$  is the WCRT<sup>4</sup> of  $\tau_s$ ).

## 4 PERIOD SELECTION

The actual periods for the security tasks are not known – we need to find the periods that ensures schedulability and gives us better monitoring frequency. Mathematically this can be expressed as the following optimization problem: minimize  $\sum_{T_s, \forall \tau_s \in \Gamma_S} T_s$ , subject to  $\mathcal{R}_s \leq T_s \leq T_s^{max}, \forall \tau_s \in \Gamma_S$ . This is a non-trivial problem since the period of  $\tau_s$  can be anything in  $[\mathcal{R}_s, T_s^{max}]$  and the response time  $\mathcal{R}_s$  is variable as it depends on the period of other higher priority security tasks. We first derive the WCRT of the security tasks and use it as a (lower) bound to find the periods. Our WCRT calculation for security tasks is based on the existing iterative analysis for global multicore scheduling [37–39] and we modify it to account the fact that RT tasks are partitioned.

#### 4.1 Preliminaries

We start by briefly reviewing the relevant terminology and parameters. We are interested in determining the response time of a job  $\tau_s^k$  of task  $\tau_s$  (*e.g.*, job under analysis) using an iterative method and the response time in each iteration is denoted by *x*.

Definition 1 (Busy Period). The busy period of  $\tau_s^k$  is the maximal continuous time interval  $[t_1, t_2)$  (until  $\tau_s^k$  finishes) where all the cores are executing either higher priority tasks or  $\tau_s^k$  itself.

Definition 2 (Interference). Given task  $\tau_i$ , the interference  $I_{\tau_s \leftarrow \tau_i}$  caused by  $\tau_i$  on  $\tau_s^k$  is the number of time units in the busy period when  $\tau_i$  executes while  $\tau_s^k$  does not.

Note that the job under analysis  $\tau_s^k$  cannot execute if all cores are busy with higher priority tasks; hence, the length of the busy period is at most  $\left\lfloor \frac{\Omega_s}{M} \right\rfloor + C_s$  by definition, where  $\Omega_s$  is the sum of the interference caused by all higher priority tasks on  $\tau_s^k$ . To compute the value of  $I_{\tau_s \leftarrow \tau_i}$ , we rely on the concept of *workload*.

Definition 3 (Workload). The workload  $W_i(x)$  of a task  $\tau_i$  in a window of length x represents the accumulated execution time of  $\tau_i$  within this time interval.

It remains to compute the workload and corresponding interference for each higher priority task  $\tau_i$ . We first show how to do so for RT tasks and then for security tasks with higher priority than  $\tau_s$ .

#### 4.2 Interference Calculation for RT Tasks

Since RT tasks are statically partitioned to cores and they have higher priority than any task that is allowed to migrate between cores, the worst-case workload for RT tasks can be trivially obtained

<sup>&</sup>lt;sup>4</sup>The calculation of WCRT is presented in Section 4.4.



Figure 2: Workload of the RT tasks for a window of size x. The arrival time of the task  $\tau_i$  is denoted by  $a_i$ .

based on the same critical instant used for single core fixed-priority scheduling case [36].

LEMMA 1. For a given core  $\pi_m$ , the maximum workload of RT tasks executed on  $\pi_m$  in any possible time interval of length x is obtained when all RT tasks are released synchronously at the beginning of the interval.

PROOF. Since RT tasks are partitioned and they have higher priorities than security tasks, the schedule of RT tasks executed on  $\pi_m$  does not depend on any other task in the system. Now consider any interval [t, t + x) of length x. We show that we can obtain an interval [t', t' + x) where all tasks are released at t', such that the workload of RT tasks on  $\pi_m$  is higher in [t', t' + x) compared to [t, t + x).

First step: let t' be the earliest time such that  $\pi_m$  continuously executes RT tasks in [t', t); if such time does not exist, then let t' = t. By definition,  $\pi_m$  does not execute RT tasks at time t' - 1. Also since RT tasks continuously execute in [t', t), the workload of RT tasks in [t', t' + x) cannot be smaller than the workload in [t, t + x).

Second step: since  $\pi_m$  is idle at t' - 1, no job of RT tasks on  $\pi_m$  released before t' can contribute to the workload in [t', t). Hence, the workload can be maximized by anticipating the release of each RT task  $\tau_r$  so that it corresponds with t'. This concludes the proof.

Let  $\Gamma_R^{\pi_m} \subseteq \Gamma_R$  denote the set of RT tasks partitioned to core  $\pi_m$ . Based on Lemma 1, an upper bound to the workload of RT tasks on  $\pi_m$  can be obtained by assuming that each RT task  $\tau_r$  is released at the beginning of the interval and each job of  $\tau_r$  executes as early as possible after being released, as shown in Fig. 2. We thus obtain the workload for RT task  $\tau_r$ :

$$W_r^R(x) = \left\lfloor \frac{x}{T_r} \right\rfloor C_r + \min(x \bmod T_r, C_r),$$
(2)

and summing over all RT tasks on  $\pi_m$  yields a total workload  $\sum_{\tau_i \in \Gamma_R^{\pi_m}} W_i^R(x)$ . Finally, we notice that by definition the interference

caused by a group of tasks executing on the same core  $\pi_m$  on  $\tau_s$  cannot be greater than  $x - C_s + 1$ . Therefore, the maximum interference caused by RT tasks on  $\pi_m$  to  $\tau_s$  can be bounded as:

$$I_{\tau_s \leftarrow \Gamma_R^{\pi_m}}\left(x, \sum_{\tau_i \in \Gamma_R^{\pi_m}} W_i^R(x)\right) = \min\left(\sum_{\tau_i \in \Gamma_R^{\pi_m}} W_i^R(x), x - C_s + 1\right).$$
(3)

The '+1' term in the upper bound of the interference (*e.g.*, Eq. (3)) ensures the convergence of iterative search for the response time (recall from Section 4.1 that at each iteration the response time is denoted by *x*) to the correct value [40]. For example, when the iterative search for the response time is started with  $x = C_s$  (*i.e.*,

Busy Period			
$a_s - t_0$	$f_s - a_s$		
t <sub>0</sub> Arri	val $(a_s)$	Finish (fs	$\rightarrow$ Ime

Figure 3: Extension of busy period for bounding the number of carry-in higher priority security tasks.



Figure 4: Illustration of carry-in task for a window of size x.

 $x - C_s = 0$ ), the search would stop immediately (and outputs an incorrect WCRT) since min  $\left(\sum_{\tau_i \in \Gamma_R^{\pi_m}} W_i^R(x), x - C_s\right) = 0.$ 

## 4.3 Interference Calculation for Security Tasks

We next consider the workload of security tasks with higher priority than  $\tau_s$ . The workload computation depends on the arrival time of the task relative to the beginning of the busy period, as specified in the following definition.

Definition 4 (Carry-in). A task  $\tau_i$  is called a *carry-in* task if there exists one job of  $\tau_i$  that has been released before the beginning of a given time window of length x and executes within the window. If no such job exists,  $\tau_i$  is referred to as a *non-carry-in* task.

Generally (but not always), the workload of a task  $\tau_i$  in the busy period is higher if  $\tau_i$  is a carry-in task than a non-carry-in task. Hence, it is important to limit the number of higher priority carry-in tasks. To this end, we follow an approach similar to prior research [37, 39] and extend the busy period of  $\tau_s^k$  from its arrival time (denoted by  $a_s$ ) to an earlier time instance  $t_0$  (see Fig. 3) such that during any time instance  $t \in [t_0, a_s)$  all cores are busy executing tasks with higher priority than  $\tau_s$ . Note that by definition, this implies that there was at least one free core (*i.e.*, not executing higher priority tasks) at time  $t_0 - 1$ .

LEMMA 2. At most M - 1 higher priority tasks can have carry-in at time  $t_0$ .

PROOF. The maximum number of higher priority tasks that can have carry-in at  $t_0$  is M - 1 since by definition there have to be strictly less than M higher priority tasks active at time  $t_0 - 1$  (otherwise they will occupy all the cores).

Since Lemma 2 holds for all tasks with higher priority than  $\tau_s$ , an immediate corollary is that the number of security tasks with carry-in at  $t_0$  also cannot be larger than M - 1. If a security task  $\tau_i$  does not have carry-in, its workload is maximized when the task is released at the beginning of the busy interval. Hence, we can calculate the workload bound  $W_i^{S_{NC}}(x)$  for the interval x using Eq. (2), e.g.,  $W_i^{S_{NC}}(x) = \left\lfloor \frac{x}{T_i} \right\rfloor C_i + \min(x \mod T_i, C_i)$ . Likewise, the workload bound for a carry-in security task  $\tau_i$  in an interval of length x starting at  $t_0$  is given by (see Fig. 4):

$$W_i^{S_{CI}}(x) = W_i^{S_{NC}} \left( \max(x - \bar{x}_i, 0) \right) + \min(x, C_i - 1), \qquad (4)$$

where  $\bar{x}_i = C_i - 1 + T_i - \mathcal{R}_i$ . We can bound the workload of the first carry-in job to  $C_i - 1$  because the job must have started executing at the latest at  $t_0 - 1$  (given that not all cores are busy). Finally, using the same argument as in Section 4.2, the interference of  $\tau_i$  can be bounded as follows:

$$I_{\tau_{s} \leftarrow \tau_{i}}(x, W_{i}(x)) = \min(W_{i}(x), x - C_{s} + 1),$$
(5)

where  $W_i(x)$  is either  $W_i^{S_{NC}}(x)$  or  $W_i^{S_{CI}}(x)$ . Notice that the WCRT and periods of security task in the carry-in workload function (see Eq. (4)) is actually an unknown parameter. However, we follow an iterative scheme that allows us to calculate the period and WCRT of all higher priority security tasks before we calculate the interference for task  $\tau_s$  (refer to Section 4.5 for details).

#### 4.4 **Response Time Analysis**

Let  $hp_S(\tau_s)$  denote the set of security tasks with a higher priority than  $\tau_s$ . Note that we do not know which (at most) M - 1 security tasks in  $hp_S(\tau_s)$  have carry-in. In order to derive the WCRT of  $\tau_s$ , let us define  $\mathcal{Z}_{\tau_s} \subset \Gamma \times \Gamma$  as the set of all partitions of  $hp_S(\tau_s)$  into two subsets  $\Gamma_s^{NC}$  and  $\Gamma_s^{CI}$  (e.g., the non overlapping set of carry-in and non-carry-in tasks) such that:

$$\Gamma_s^{NC} \cap \Gamma_s^{CI} = \emptyset, \Gamma_s^{NC} \cup \Gamma_s^{CI} = hp_S(\tau_s), \text{ and } |\Gamma_s^{CI}| \le M - 1,$$

*e.g.*, there are at most M - 1 carry-in tasks.

For a given carry-in and non-carry-in set (e.g.,  $\Gamma_s^{NC}$  and  $\Gamma_s^{CI}$ ), we can calculate the total interference experienced by  $\tau_s$  as follows:

$$\Omega_{s}(x,\Gamma_{s}^{NC},\Gamma_{s}^{CI}) = \sum_{\pi_{m}\in\mathcal{M}} I_{\tau_{s}\leftarrow\Gamma_{R}^{\pi_{m}}}\left(x,\sum_{\tau_{i}\in\Gamma_{R}^{\pi_{m}}}W_{i}^{R}(x)\right) + \sum_{\tau_{i}\in\Gamma_{s}^{NC}} I_{\tau_{s}\leftarrow\tau_{i}}\left(x,W_{i}^{S_{NC}}(x)\right) + \sum_{\tau_{i}\in\Gamma_{s}^{CI}} I_{\tau_{s}\leftarrow\tau_{i}}\left(x,W_{i}^{S_{CI}}(x)\right).$$
(6)

For a given  $\Gamma_s^{NC}, \Gamma_s^{CI}$  sets response time  $\mathcal{R}_{s|(\Gamma_s^{NC}, \Gamma_s^{CI})}$  will be the minimal solution of the following iteration<sup>5</sup> [37]:

$$x = \left\lfloor \frac{\Omega_s(x, \Gamma_s^{NC}, \Gamma_s^{CI})}{M} \right\rfloor + C_s.$$
(7)

We can solve this using an iterative fixed-point search with the initial condition  $x^{(0)} = C_s$ . The search terminates if there exists a solution (*i.e.*,  $x = x^{(k)} = x^{(k-1)}$  for some iteration k) or when  $x^{(k)} >$  $T_s^{max}$  for any iteration k since  $\tau_s$  becomes trivially unschedulable for WCRT greater than  $T_s^{max}$ . Finally we can calculate the WCRT of  $\tau_s$  as follows:

$$\mathcal{R}_{s} = \max_{\left(\Gamma_{s}^{NC}, \Gamma_{s}^{CI}\right) \in \mathcal{Z}_{\tau_{s}}} \mathcal{R}_{s \mid (\Gamma_{s}^{NC}, \Gamma_{s}^{CI})}.$$
(8)

# 4.5 Algorithm

The security task  $\tau_s$  remains schedulable with any period  $T_s \in$  $[\mathcal{R}_s, T_s^{max}]$ . However as mentioned earlier, the calculation of  $\mathcal{R}_s$ requires us to know the period and response time of other high priority tasks  $\tau_h \in hp_S(\tau_s)$ . Also if we arbitrarily set  $T_s = \mathcal{R}_s$  (since this allows us to execute security tasks more frequently) it may negatively affect the schedulability of other tasks that are at a lower priority than  $\tau_s$  because of a high degree of interference from  $\tau_s$ .

#### Algorithm 1 Period Selection

**Input:** Set of real-time and security tasks  $\Gamma = \Gamma_R \cup \Gamma_S$ 

- Output: Periods of the security tasks, T (if the security tasks are schedulable); Unschedulable otherwise
- 1: Set  $T_s := T_s^{max}$  and calculate  $\mathcal{R}_s$  for  $\forall \tau_s \in \Gamma_S$
- if  $\exists \tau_s$  such that  $\mathcal{R}_s > T_s^{max}$  then 2:
- 3: return Unschedulable
- 4: end if
- 5: for each security task  $\tau_s \in \Gamma_S$  (from higher to lower priority) do
- 6: 7:
- /\* Find period for which all lower priority tasks are schedulable \*/ Find minimum  $T_s^* \in [\mathcal{R}_s, T_s^{max}]$  using Algorithm 2 such that  $\forall \tau_j \in lp(\tau_s)$  remains schedulable (e.g.,  $\mathcal{R}_j \leq T_j^{max})$ )
- 8: Update  $\mathcal{R}_i$  for  $\forall \tau_i \in lp(\tau_s)$  considering the interference with new period  $T_s^*$ 9: end for

10: return T :=  $[T_s^*]_{\forall \tau_s \in \Gamma_S}$  /\* return the periods \*/

Hence, we developed an iterative algorithm that gives us a trade-off between schedulability and monitoring frequency.

Our proposed solution (refer to Algorithm 1 for a formal description) works as follows. We first fix the period of each security task  $T_s^{max}$  and calculate the response time  $\mathcal{R}_s$  using the approach presented in Section 4.3 (Line 1). If there exists a task  $\tau_j$  such that  $\mathcal{R}_j > T_i^{max}$  we report the taskset as unschedulable (Line 2) since it is not possible to find a period for the security tasks within the designer provided bounds - this unschedulability result will help the designer in modifying the requirements (and perhaps RT tasks' parameters, if possible) accordingly to integrate security tasks for the target system. If the taskset is schedulable with  $T_s^{max}$ , we then iteratively optimize the periods from higher to lower priority order (Lines 5-9) and return the period (Line 10). To be specific, for each task  $\tau_s \in \Gamma_S$  we perform a logarithmic search [41, Ch. 6] (see Algorithm 2 for the pseudocode) and find the minimum period  $T_s^*$ within the range  $[R_s, T_s^{max}]$  such that all low priority tasks (denoted as  $lp(\tau_s)$ ) remain schedulable, *e.g.*,  $\forall \tau_j \in lp(\tau_s) : \mathcal{R}_j \leq T_j^{max}$  (Line 7). Note that since we perform these steps from higher to lower priority order, WCRT and period of all higher priority tasks (e.g.,  $\forall \tau_h \in hp(\tau_s)$ ) are already known. We then update the response times of all low priority task  $\tau_i \in lp(\tau_s)$  considering the interference from the newly calculated period  $T_s^*$  (Line 8) and repeat the search for next security task.

#### 5 **EVALUATION**

We evaluate HYDRA-C on two fronts: (i) a proof-of-concept implementation on an ARM-based rover platform with security applications - to demonstrate the viability of our scheme in a realistic setup (Section 5.1); and (ii) with synthetically generated workloads for broader design-space exploration (Section 5.2). Our implementation code will be made available in a public, opensourced repository [42].

#### **Experiment with an Embedded Platform** 5.1 and Security Applications

5.1.1 Platform Overview. We implemented our ideas on a rover platform manufactured by Waveshare [43]. The rover hardware/peripherals (e.g., wheel, motor, servo, sensor, etc.) are controlled by a Raspberry Pi 3 (RPi3) Model B [44] SBC (single board computer). The RPi3 is equipped with a 1.2 GHz 64-bit quad-core ARM Cortex-A53 CPU on top of Broadcom BCM2837

<sup>&</sup>lt;sup>5</sup>Note that the worst-case is when the job arrives at  $t_0$  (*i.e.*,  $a_s = t_0$ ).

# **Algorithm 2** Calculation of Minimum Feasible Period for the Security Task $\tau_s$

**Input:** Set of real-time and security tasks  $\Gamma = \Gamma_R \cup \Gamma_S$ **Output:** A feasible period  $T_s^*$  for the security task under analysis (*i.e.*,  $\tau_s$ ) 1: Define  $T_s^l := \mathcal{R}_s, T_s^r := T_s^{max}, T_s^c := 0$ 2: Set  $\widehat{\mathcal{T}}_s := \{T_s^{max}\} / *$  Initialize a variable to store the set of feasible periods \*/ while  $T_s^l <= T_s^r$  do 3: Update  $T_s^c := \lfloor \frac{T_s^l + T_s^r}{2} \rfloor$ 4: if  $\exists \tau_j \in lp(\tau_s)$  such that  $\tau_j$  is not schedulable with  $T_s = T_s^c$  then 5: Increase the period of  $\tau_s$  to make the taskset schedulable (e.g., by reducing 6: the interference) \*/ 7: Update  $T_s^l := T_s^c + 1$ 8: else Taskset is schedulable with  $T_c^c * /$ 9:  $\widehat{\mathcal{T}}_s := \widehat{\mathcal{T}}_s \cup \{T_s^c\} / Add T_s^c$  to the feasible period list \*/ 10: 11: Check schedulability with smaller period for next iteration \*/ 12: Update  $T_s^r := T_s^c - 1$ 13: end if 14: end while 15: Set  $T_s^* := \min(\widehat{\mathcal{T}}_s)$  /\* Find the minimum period from the set of feasible periods \*/ 16: return  $T_s^*$  /\* return the period of  $\tau_s^*$  //

SoC (system-on-chip). In our experiments we focus on a dualcore setup (*e.g.*, activated only core0 and core1) and disabled the other two cores) – this was done by modifying the boot command file /boot/cmdline.txt and set the flag maxcpus=2. The base hardware unit of the rover is connected with RPi3 using a 40-pin GPIO (general-purpose input/output) header. The rover supports omni-directional movement and can move avoiding obstacles using an infrared sensor (*e.g.*, ST188 [45]). We also attached a camera (RPi3 camera module) that can capture static images (3280 × 2464 pixel resolution). The detailed specifications of the rover hardware (*e.g.*, base chassis, adapter, *etc.*) are available on the vendor website [43].

5.1.2 Experiment Setup and Implementation. We implemented our security integration scheme in Linux kernel 4.9 and enabled real-time capabilities by applying the PREEMPT\_RT patch [34] (version 4.9.80-rt62-v7+). In our experiments the rover moved around autonomously and periodically captured images (and stored them in the internal storage). We assumed implicit deadlines for RT tasks and considered two RT tasks: (a) a navigation task – that avoids obstacles (by reading measurements from infrared sensor) and navigates the rover and (b) a camera task that captures and stores still images. We do not make any modifications to the vendor provided control codes (e.g., navigation task). In our experiments we used the following parameters  $(C_r, T_r)$ : (240, 500) ms and (1120, 5000) ms, for navigation and camera tasks, respectively (i.e., total RT task utilization was 0.7040). We calculated the WCET values using ARM cycle counter registers (CCNT) and set periods in a way that the rover can navigate and capture images without overloading the RPi3 CPU. Since CCNT is not accessible by default, we developed a Linux loadable kernel module and activated the registers so that our measurement scripts can access counter values.

To integrate security into this rover platform, we included two additional security tasks: (*a*) an open-source security application, Tripwire [17], that checks intrusions in the image data-store and (*b*) our custom security task that checks current kernel modules (as a preventive measure to detect rootkits) and compares with an expected profile of modules. The WCET of the security tasks were 5342 ms and 223 ms, respectively and the maximum periods of

**Table 2: Summary of the Evaluation Platform** 

Artifact	Configuration/Tools
Platform	1.2 GHz 64-bit Broadcom BCM2837
CPU	ARM Cortex-A53
Memory	1 Gigabyte
Operating System	Debian Linux (Raspbian Stretch Lite)
Kernel version	Linux Kernel 4.9
Real-time patch	PREEMPT_RT 4.9.80-rt62-v7+
Kernel flags	CONFIG_PREEMPT_RT_FULL enabled
Boot parameters	<pre>maxcpus=2, force_turbo=1,</pre>
	arm_freq=700, arm_freq_min=700
WCET measurement	ARM cycle counter registers
Task partition	Linux taskset



Figure 5: Experiments with rover platform: (a) time (cycle counts) to detect intrusions; (b) average number of context switches. On average our scheme can detect the intrusions faster without impacting the performance of RT tasks.

security tasks were assumed to be 10000 ms (*e.g.*, total system utilization is at least 0.7040 + 0.5565 = 1.2605) – we picked this maximum period value by trial and error so that the taskset became schedulable for demonstration purposes. We used the Linux taskset utility [46] for partitioning tasks to the cores and the tasks were scheduled using Linux native sched\_setscheduler() function. For accuracy of our measurements we disabled all CPU frequency scaling features in the kernel and executed RPi with a constant frequency (*e.g.*, 700 MHz – the default value). The system configurations and tools used in our experiments are summarized in Table 2.

We compared the performance of our scheme with prior work, HYDRA [26]. In that work, researchers proposed to statically partition the security tasks among the multiple cores – to our knowledge HYDRA is the state-of-the-art mechanism for integrating security in legacy multicore-based RT platforms. The key idea in HYDRA was to allocate security tasks using a greedy best-fit strategy: for each task, allocate it to a core that gives maximum monitoring frequency (*i.e.*, shorter period) without violating schedulability constraints of already allocated tasks.

5.1.3 Experience and Evaluation. We observed the performance of HYDRA-C by analyzing how quickly an intrusion can be detected. We considered the following two realistic attacks<sup>6</sup>: (*i*) an ARM

 $<sup>^6\</sup>underline{\text{Note:}}$  our focus here is on the integration of any given security mechanisms rather the detection of any particular class of intrusions. Hence we assumed that there were no zero-day attacks and the security tasks were able the detect the corresponding attacks correctly (*i.e.*, there were no false-positive/negative errors) – although the

**Table 3: Simulation Parameters** 

Parameter	Values
Process cores, M	{2,4}
Number of real-time tasks, $N_R$	$[3 \times M, 10 \times M]$
Number of security tasks, $N_S$	$[2 \times M, 5 \times M]$
Period distribution (RT and security tasks)	Log-uniform
RT task allocation	Best-fit
RT task period, $T_r$	[10, 1000] ms
Maximum period for security tasks, $T_s^{max}$	[1500, 3000] ms
Minimum utilization of security tasks	At least 30% of
	RT tasks
Base utilization groups	10
Number of taskset in each configuration	250

shellcode [47] that allows the attacker to modify the contents of the image data-store – this attack can be detected by Tripwire; (*ii*) a rootkit [48] that intercepts all the read() system calls – our custom security task can detect the presence of the malicious kernel module that is used to inject the rootkit. For each of our experimental trials we launched attacks at random points during program execution (*i.e.*, from the RT tasks) and used ARM cycle counters to measure the detection time. In Fig. 5a we show the average time to detect both the intrusions (in terms of cycle counts, collected from 35 trials) for HYDRA-C and HYDRA schemes. From our experiments we found that, on average, our scheme can detect intrusions 19.05% faster compared to the HYDRA approach (Fig. 5a). Since our scheme allows security tasks to migrate across cores, it provides smaller response time (*e.g.*, shorter period) in general and that leads to faster detection times.

We next measured the overhead of our security integration approach in terms of number of context switches (CS). For each of the trials we observed the schedule of the RT and security tasks for 45 seconds and counted the number of CS using the Linux perf tool [22]. In Fig. 5b we show the number of CS (y-axis in the figure) for HYDRA-C and HYDRA schemes (for 35 trials). As shown in the figure, our approach increases the number of CS (since we permit migration across cores) compared to the other scheme that statically partitions security tasks. From our experiments we found that, on average, our scheme increases CS by 1.75 times. However, this increased CS overhead *does not impact the deadlines of RT tasks* (since the security tasks always execute with a priority lower than the RT tasks) and thus may be acceptable for many RT applications.

### 5.2 Experiment with Synthetic Tasksets

We also conducted experiments with (randomly generated) synthetic workloads for broader design-space exploration.

5.2.1 Taskset Generation and Parameters. In our experiments we used parameters similar to those in related work [15, 25, 26, 38, 49, 50] (see Table 3). We considered  $M \in \{2, 4\}$  cores and each taskset instance contained  $[3 \times M, 10 \times M]$  RT and  $[2 \times M, 5 \times M]$  security tasks. To generate tasksets with an even distribution of tasks, we grouped the real-time and security tasksets by base-utilization from [(0.01 + 0.1i)M, (0.1 + 0.1i)M] where  $i \in \mathbb{Z}, 0 \leq 1$ 



Figure 6: Euclidean distance between achievable period and maximum period vectors for different utilizations. Larger distance (y-axis in the figure) implies security tasks execute more frequently.

 $i \leq 9$ . Each utilization group contained 250 tasksets (*e.g.*, total  $10 \times 250 = 2500$  tasksets were tested for each core configuration). We only considered the schedulable tasksets (*e.g.*, the condition in Section 2.1 was satisfied for all RT tasks) – since tasksets that fail to meet this condition are trivially unschedulable. Task periods were generated according to a log-uniform distribution. Each RT task had periods between [10, 1000] ms and the maximum periods for security tasks were selected from [1500, 3000] ms. We assumed that RT tasks were partitioned using a best-fit [1] strategy and the total utilization of the security tasks was at least 30% of the system utilization. For a given number of tasks and total system utilization, the utilization of individual tasks were generated using Randfixedsum algorithm [51].

5.2.2 Impact on Inter-Monitoring Interval. We first observe how frequently we can execute (schedule) security tasks compared to the designer specified bound (Fig. 6). The x-axis of Fig. 6 shows the normalized utilization  $\frac{U}{M}$  where U is the minimum utilization requirement and given as follows:  $U = \sum_{T_r \in \Gamma_R} \frac{C_r}{T_r} + \sum_{\tau_s \in \Gamma_S} \frac{C_s}{T_s^{max}}$ . The y-axis represents the Euclidean distance between the calculated period vector  $\mathbf{T}^* = [T_s^*]_{\forall \tau_s \in \Gamma_S}$  and maximum period vector  $\mathbf{T}^{max} = [T_s^{max}]_{\forall \tau_s \in \Gamma_S}$  (normalized to 1). A higher distance implies that tasks can run more frequently. As we can see from the figure for higher utilizations, the distance reduces (*e.g.*, periods are closer to the maximum value) – this is mainly due to the interference from higher priority (RT and security) tasks. The results from this figure suggest that we can execute security tasks more frequently for low to medium utilization cases.

5.2.3 Impact on Schedulability and Security Trade-off. While in this work we consider a legacy RT system (*i.e.*, where RT tasks are partitioned to respective cores), for comparison purposes we considered the following two schemes (in addition to the related work, HYDRA, introduced in Section 5.1) that do not consider any period adaptation for security tasks.

• GLOBAL-TMax: In this scheme all the RT and security tasks are scheduled using a global fixed-priority multicore scheduling scheme [30]. Since our focus here is on schedulability we set  $T_s = T_s^{max}$ ,  $\forall \tau_s \in \Gamma_S$  (recall that a taskset can be considered schedulable if the following conditions hold:  $\mathcal{R}_r \leq D_r$ ,  $\forall \tau_r \in \Gamma_R$  and  $\mathcal{R}_s \leq T_s^{max}$ ,  $\forall \tau_s \in \Gamma_S$ ). This scheme allows us to observe the performance

generic framework proposed in this paper allows the designers to accommodate any desired security (*e.g.*, intrusion detection/prevention) technique.



Figure 7: Impact on schedulability and security. (*a*) The acceptance ratio vs taskset utilizations for 2 and 4 core platforms: our scheme outperforms HYDRA and GLOBAL-TMax approaches for higher utilizations. (*b*) Difference in period vectors for our approach and reference schemes (*e.g.*, HYDRA, GLOBAL-TMax, HYDRA-TMax): the non-negative distance (y-axis in the figure) implies that HYDRA-C finds shorter periods than other schemes.

impacts of binding RT tasks to the cores (due to legacy compatibility).

• HYDRA-TMax: This is similar to the HYDRA approach introduced in Section 5.1 (*i.e.*, security tasks were partitioned using best-fit allocation as before) but instead of minimizing periods here we set  $T_s = T_s^{max}, \forall \tau_s$ . This allows us to observe the trade-offs between schedulability and security in a fully-partitioned system.

In Fig. 7a we compare the performance of HYDRA-C with the HYDRA, GLOBAL-TMax and HYDRA-TMax strategies in terms of acceptance ratio (y-axis in the figure) defined as the number of schedulable tasksets (*e.g.*,  $\mathcal{R}_s \leq T_s^{max}, \forall \tau_s$ ) over the generated one and the x-axis shows the normalized utilization  $\frac{U}{M}$ . As we can see from the figure, HYDRA-C outperforms HYDRA when the utilization increases (*i.e.*,  $\frac{U}{M} > 0.2$ ). This is because our scheme allows security tasks to execute in parallel across cores and also allocate periods considering the schedulability constrains of all low priority tasks - this results in a smaller response time and can find more tasksets that satisfy the designer specified bound. In contrast HYDRA uses a greedy approach that minimizes the periods of higher priority tasks first without considering the global state. Also HYDRA statically binds the security task to the core and hence suffers interference from the higher priority tasks assigned to that core - this leads to lower acceptance ratios. For higher utilizations (*i.e.*,  $\frac{U}{M} \ge 0.7$ ) HYDRA-C can find tasksets schedulable that can not be easily partitioned by using the HYDRA-TMax scheme. The acceptance ratio of our method and the HYDRA-TMax scheme is equal when  $\frac{U}{M} < 0.7$ . This is because, for lower utilizations some lower priority security tasks experience less interference due to longer periods and specific core assignment (recall we set  $T_s = T_s^{max}$  for all security tasks). While we bind the RT tasks to cores (due to legacy compatibility), it does not affect the schedulability (i.e., the acceptance ratio of HYDRA-C is higher when compared to the GLOBAL-TMax scheme). This is because, RT tasks are already schedulable when partitioned (e.g., by assumption on taskset generation, see Section 5.2.1) and our analysis reduces the interference that RT tasks have on security ones. For higher utilizations, the acceptance ratio drops for all the schemes since it is not possible to satisfy all the constraints due to

the high interference from RT and security tasks. We also highlight that while our approach results in better schedulability, HYDRA-C/HYDRA-TMax (*i.e.*, where legacy RT tasks are partitioned to the cores) and GLOBAL-TMax (*i.e.*, where all tasks can migrate) schemes are incomparable in general (*e.g.*, there exists taskset that may be schedulable by task partitioning but not in global scheme where migration is allowed and vice-versa) – we allow security tasks to migrate due to security requirements (*e.g.*, to achieve faster intrusion detection – as we explain in the next experiments, see Fig. 7b).

In the final set of experiments (Fig. 7b) we compare the achievable periods (in terms of Euclidean distance) for our approach and the other schemes. The x-axis in the Fig. 7b shows the normalized utilizations and the y-axis represents the average difference between the following period vectors: (a) between HYDRA-C and HYDRA (dashed line); (b) HYDRA-C and other strategies (e.g., GLOBAL-TMax and HYDRA-TMax) that do not consider period minimization (dotted marker) for dual and quad core setup. Higher distance values imply that the periods calculated by HYDRA-C are smaller (i.e., leads to faster detection time) and our approach outperforms the other scheme. For low to medium utilizations (e.g.,  $0.2 \le U \le 0.5$ ) HYDRA-C performs better when compared to HYDRA. In situations with higher utilizations, the lesser availability of slack time results in HYDRA-C and HYDRA performing in a similar manner. Also, for higher utilizations HYDRA is unable to find schedulable tasksets and hence there exist fewer data points.

Our experiments also show that compared to GLOBAL-TMax and HYDRA-TMax our approach finds smaller periods in most cases (Fig. 7b). This is expected since there is no period adaptation (*i.e.*, we set  $T_s = T_s^{max}$  for those schemes). However it is important to note that HYDRA-C achieves better execution frequency (*i.e.*, smaller periods) without sacrificing schedulability as seen in Fig. 7a. That is, our semi-partitioned approach achieves better continuous monitoring when compared with both a fully-partitioned approach (HYDRA, HYDRA-TMax) and a global scheduling approach (GLOBAL-TMax) while providing the same or better schedulability.

# 6 **DISCUSSION**

In this paper we do not design for any specific security tasks (the IDS system used is meant for demonstration purposes only) and allow designers to integrate their preferred techniques. Depending on the actual implementation of the security tasks some attack may not be detectable. For instance, the system may be vulnerable to zero-day attacks if the security tasks are not designed to detect unforeseen exploits or anomalous behaviors. There exists cases where security tasks may require some amount of system modifications and/or porting efforts – say a timing behavior based security checking [11, 28, 52] may require the insertion of probing mechanisms inside the RT application tasks (or additional hardware) so that security tasks can validate their expected execution profiles.

HYDRA-C abstracts security tasks (and underlying monitoring events) and works in a proactive manner. However, designers may want to integrate security tasks that *react*, based on anomalous behavior. For instance, let at time *t*, *j*-th job of task  $\tau_s$  (*e.g.*,  $\tau_s^j$ ) performs action a<sub>0</sub> (*e.g.*, runtime of real-time tasks). Because of intrusions (or perhaps due to other system artifacts) in time [*t*, *t*+*T*<sub>s</sub>] (*T*<sub>s</sub> is the period of  $\tau_s$ ), job  $\tau_s^{j+1}$  finds that a<sub>0</sub> is not behaving as expected. Therefore  $\tau_s^{j+1}$  may perform both actions, a<sub>0</sub> and a<sub>1</sub> (say that checks the list of system calls, to see if any undesired calls are executed). One way to support such a feature is to consider the dependency (*i.e.*, a<sub>1</sub> depends on a<sub>0</sub> in this case) between security checks (*e.g.*, sub-tasks). We intend to extend our framework considering dependency between security tasks.

#### 7 RELATED WORK

RT Scheduling and Period Optimization. Although not in the context of RT security, the scheduling approaches present in this paper can be considered as a special case of prior work [53] where each task can bind to any arbitrary number of available cores. For a given period, this prior analysis [53] is pessimistic for the model considered by HYDRA-C (i.e., RT tasks are partitioned and security tasks can migrate on any core) in a sense that it overapproximates carry-in interference from the tasks bound to single cores (e.g., RT tasks) and hence results in lower schedulability (e.g., identical to the GLOBAL-TMax scheme in Fig. 7a). Researchers also propose various semi-partitioned scheduling strategies for fixedpriority RTS [35, 54]. However, this work primarily focuses on improving schedulability (e.g., by allowing highest priority task to migrate) and they are not designed for security requirements in consideration (e.g., minimizing periods and executing security tasks with fewer interruption for faster anomaly detection). There exists other work [55] in which the authors statically assign the periods for multiple independent control tasks considering control delay as a cost metric. Davare et al. [56] propose to assign task and message periods as well as satisfy end-to-end latency constraints for distributed automotive systems. While the previous work focus on optimizing period of all the tasks in the system for a single core setup, our goal is to ensure security without violating timing constraints of the RT tasks in a multicore platform.

Security Solutions for RTS. In recent years researchers proposed various mechanisms to provide security guarantees into legacy and non-legacy RTS (both single and multicore platforms) in several directions, *viz.*, integration of security mechanisms [25–27], authenticating/encrypting communication channels [7–10, 14, 57], side-channel defence techniques [15, 58–61] as well as hardware/software-based frameworks [11–13, 62–65].

Perhaps the closest line of research is HYDRA [26] where authors proposed to statically partition security tasks to the cores and used an optimization-based solution to obtain the periods. While this approach does not have the overhead of context switches across cores, as we observed from our experiments (Section 5), that scheme results in a poor acceptance ratio for larger utilizations, and suffers interference from other high priority tasks leading to slower detection of intrusions (i.e., less effective). The problem of integrating security for single core RTS is addressed in prior research [25] where authors used hierarchical scheduling [66] and proposed to execute security tasks with a low priority server. This approach is also extended to a multi-mode framework [27] that allows security tasks to execute in different modes (i.e., passive monitoring with lowest priority as well as exhaustive checking with higher priority). These server-based approaches, however, may require additional porting efforts for legacy systems.

There exists recent work [9, 10] to secure cyber-physical systems from man-in-the-middle attacks by enabling authentication mechanisms and timing-aware network resource scheduling. There has also been work [7, 8, 14] where authors proposed to add protective security mechanisms into RTS and considered periodic task scheduling where each task requires a security service whose overhead varies according to the required level of service. The problem of designing secure multi-mode RTS have also been addressed in prior work [57] under dynamic-priority scheduling. In contrast, we consider a multicore fixed-priority scheduling mechanism where security tasks are executed periodically, across cores, while meeting real-time requirements. The above mentioned work are designed for single core platforms and it is not straightforward to retrofit those approaches for multicore legacy systems.

In another direction, the issues related to information leakage through storage timing channels using shared architectural resources (e.g., caches) is introduced in prior work [15, 58, 59]. The key idea is to use a modified fixed-priority scheduling algorithm with a state cleanup mechanism to mitigate information leakage through shared resources. However, this leakage prevention comes at a cost of reduced schedulability. Researchers also proposed to limit inferability of deterministic RT schedulers by randomizing the task execution patterns. Yoon et al. [60] proposed a schedule obfuscation method for fixed-priority RM systems. A combined online/offline randomization scheme [61] is also proposed to reduce determinism for time-triggered (TT) systems where tasks are executed based on a pre-computed, offline, slot-based schedule. We highlight that all the aforementioned work either requires modification to the scheduler or RT task parameters, and is designed for single core systems only.

Unlike our approach that works at the scheduler-level, researchers also proposed hardware/software-based architectural solutions [11–13, 62–65, 67] to improve the security posture of future RTS. Those solutions require system-level modifications and are not suitable for legacy systems. To our knowledge this

is the first work that aims to achieve continuous monitoring for multicore-based legacy RT platforms.

#### **CONCLUSION** 8

Threats to safety-critical RTS are growing and there is a need for developing layered defense mechanisms to secure such critical systems. We present algorithms to integrate continuous security monitoring for legacy multicore-based RTS. By using our framework, systems engineers can improve the security posture of RTS. This additional security guarantee also enhances safety which is the main goal for such systems.

#### REFERENCES

- [1] J. Chen, "Partitioned multiprocessor fixed-priority scheduling of sporadic realtime tasks," in Euromicro ECRTS, 2016, pp. 251-261.
- [2] N. Falliere, L. O. Murchu, and E. Chien, "W32, stuxnet dossier," White paper, Symantec Corp., Security Response, vol. 5, p. 6, 2011.
- [3] R. M. Lee, M. J. Assante, and T. Conway, "Analysis of the cyber attack on the ukrainian power grid," SANS Industrial Control Systems, 2016.
- [4] K. Koscher, A. Czeskis, F. Roesner, S. Patel, T. Kohno, S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham et al., "Experimental security analysis of a modern automobile," in IEEE S&P, 2010, pp. 447–462.
- [5] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, T. Kohno et al., "Comprehensive experimental analyses of automotive attack surfaces," in USENIX Sec. Symp., 2011.
- [6] S. S. Clark and K. Fu, "Recent results in computer security for medical devices," in MobiHealth, 2011, pp. 111-118.
- [7] T. Xie and X. Qin, "Improving security for periodic tasks in embedded systems through scheduling," ACM TECS, vol. 6, no. 3, p. 20, 2007.
- [8] M. Lin, L. Xu, L. T. Yang, X. Qin, N. Zheng, Z. Wu, and M. Qiu, "Static security optimization for real-time systems," IEEE Trans. on Indust. Info., vol. 5, no. 1, pp. 22-37, 2009.
- [9] V. Lesi, I. Jovanov, and M. Pajic, "Network scheduling for secure cyber-physical systems," in IEEE RTSS, 2017, pp. 45–55.
- [10] V. Lesi, I. Jovanov, and M. Pajic, "Security-aware scheduling of embedded control tasks," ACM TECS, vol. 16, pp. 188:1-188:21, 2017.
- [11] M.-K. Yoon, S. Mohan, J. Choi, J.-E. Kim, and L. Sha, "SecureCore: A multicorebased intrusion detection architecture for real-time embedded systems," in IEEE RTAS, 2013, pp. 21-32.
- [12] M.-K. Yoon, S. Mohan, J. Choi, and L. Sha, "Memory heat map: anomaly detection in real-time embedded systems using memory behavior," in ACM/EDAC/IEEE DAC, 2015, pp. 1-6.
- [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 ACM/IEEE IoTDI, 2017, pp. 191-196.
- [14] T. Xie, A. Sung, and X. Qin, "Dynamic task scheduling with security awareness in real-time systems," in Parallel and Distributed Processing Symposium, 2005. Proceedings. 19th IEEE International. IEEE, 2005, pp. 8-pp.
- [15] S. Mohan, M.-K. Yoon, R. Pellizzoni, and R. B. Bobba, "Integrating security constraints into fixed priority real-time schedulers," RTS Journal, vol. 52, no. 5, pp. 644-674, 2016.
- [16] S. Mohan, "Worst-case execution time analysis of security policies for deeply embedded real-time systems," ACM SIGBED Review, vol. 5, no. 1, p. 8, 2008.
- [17] "Tripwire," https://github.com/Tripwire/tripwire-open-source. [18] "Advanced Intrusion Detection Environment (AIDE)," http://aide.sourceforge.
- net/.
- [19] "The Bro network security monitor," https://www.bro.org.
- [20] M. Roesch, "Snort lightweight intrusion detection for networks," in USENIX Conf. on Sys. Admin., 1999, pp. 229-238.
- [21] L. L. Woo, M. Zwolinski, and B. Halak, "Early detection of system-level anomalous behaviour using hardware performance counters," in DATE, 2018, pp. 485-490.
- [22] V. M. Weaver, "Linux perf\_event features and overhead," in IEEE FastPath, 2013. [23] "OProfile," http://oprofile.sourceforge.net/.
- [24] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," ACM CSUR, vol. 41, no. 3, p. 15, 2009.
- [25] M. Hasan, S. Mohan, R. B. Bobba, and R. Pellizzoni, "Exploring opportunistic execution for integrating security into legacy hard real-time systems," in IEEE RTSS, 2016, pp. 123-134.
- [26] M. Hasan, S. Mohan, R. Pellizzoni, and R. B. Bobba, "A design-space exploration for allocating security tasks in multicore real-time systems," in DATE, 2018, pp. 225 - 230.
- [27] M. Hasan, S. Mohan, R. Pellizzoni, and R. B. Bobba, "Contego: An adaptive

framework for integrating security tasks in real-time systems," in Euromicro ECRTS, 2017, pp. 23:1-23:22

- [28] M. Hamad, Z. A. Hammadeh, S. Saidi, V. Prevelakis, and R. Ernst, "Prediction of abnormal temporal behavior in real-time systems," in ACM SAC, 2018, pp. 359-367
- [29] L. Sha, M. Caccamo, R. Mancuso, J.-E. Kim, M.-K. Yoon, R. Pellizzoni, H. Yun, R. B. Kegley, D. R. Perlman, G. Arundale et al., "Real-time computing on multicore processors," IEEE Comp., vol. 49, no. 9, pp. 69-77, 2016.
- [30] R. I. Davis and A. Burns, "A survey of hard real-time scheduling for multiprocessor systems," ACM Comput. Surv., vol. 43, no. 4, pp. 35:1-35:44, 2011.
- [31] C.-Y. Chen, M. Hasan, and S. Mohan, "Securing real-time Internet-of-things," Sensors, vol. 18, no. 12, 2018.
- G. Heiser and B. Leslie, "The okl4 microvisor: Convergence point of microkernels [32] and hypervisors," in ACM APSys. ACM, 2010, pp. 19-24
- [33] F. Kolnick, "The qnx 4 real-time operating system," Basis Comp. Sys. Inc., 1998.
- [34] L. Fu and R. Schwebel, "Real-time Linux wiki," https://rt.wiki.kernel.org/index. php/rt\_preempt\_howto, [Online].
- S. Kato and N. Yamasaki, "Semi-partitioned fixed-priority scheduling on [35] multiprocessors," in IEEE RTAS, 2009, pp. 23-32.
- C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a [36] hard-real-time environment," JACM, vol. 20, no. 1, pp. 46-61, 1973
- [37] N. Guan, M. Stigge, W. Yi, and G. Yu, "New response time bounds for fixed priority multiprocessor scheduling," in IEEE RTSS, 2009, pp. 387-397.
- Y. Sun, G. Lipari, N. Guan, and W. Yi, "Improving the response time analysis of [38] global fixed-priority multiprocessor scheduling," in IEEE RTCSA, 2014, pp. 1-9.
- [39] S. Baruah, "Techniques for multiprocessor global schedulability analysis," in IEEE RTSS, 2007, pp. 119-128.
- [40] M. Bertogna and M. Cirinei, "Response-time analysis for globally scheduled symmetric multiprocessor platforms," in *IEEE RTSS*, 2007, pp. 149–160.
- D. E. Knuth, The art of computer programming: sorting and searching, 1997, vol. 3. [41] [42] "Implementation codes of the security integration framework," https://github. com/mnwrhsn/multicore-continuous-security-monitoring.
- [43] "Alphabot2 wiki," https://www.waveshare.com/wiki/AlphaBot2-Pi.
- [44] "Raspberry Pi," https://www.raspberrypi.org/products/raspberry-pi-3-model-b/.
- "ST188 Datasheet," http://www.npnec.com/en\_pdf/ST188-EN.pdf. [45]
- [46]
- "taskset(1) Linux man page," https://linux.die.net/man/1/taskset. "Linux ARM shellcode," https://www.exploit-db.com/exploits/21253/. [47]
- "Linux rootkit," https://github.com/crudbug/simple-rootkit. [48]
- [49] R. I. Davis, A. Burns, J. Marinho, V. Nelis, S. M. Petters, and M. Bertogna, "Global and partitioned multiprocessor fixed priority scheduling with deferred preemption," ACM TECS, vol. 14, no. 3, pp. 47:1-47:28, 2015.
- [50] M.-K. Yoon, J.-E. Kim, R. Bradford, and L. Sha, "Holistic design parameter optimization of multiple periodic resources in hierarchical scheduling," in DATE, 2013, pp. 1313-1318.
- [51] P. Emberson, R. Stafford, and R. I. Davis, "Techniques for the synthesis of multiprocessor tasksets," in WATERS, 2010, pp. 6-11.
- [52] M.-K. Yoon, M. Christodorescu, L. Sha, and S. Mohan, "The dragonbeam framework: Hardware-protected security modules for in-place intrusion detection," in ACM SYSTOR, 2016, pp. 1:1-1:12.
- [53] A. Gujarati, F. Cerqueira, and B. B. Brandenburg, "Schedulability analysis of the Linux push and pull scheduler with arbitrary processor affinities," in Euromicro ECRTS, 2013, pp. 69-79.
- [54] K. Lakshmanan, R. Rajkumar, and J. Lehoczky, "Partitioned fixed-priority preemptive scheduling for multi-core processors," in Euromicro ECRTS, 2009, pp. 239-248.
- [55] E. Bini and A. Cervin, "Delay-aware period assignment in control systems," in IEEE RTSS, 2008, pp. 291-300.
- A. Davare, Q. Zhu, M. Di Natale, C. Pinello, S. Kanajan, and A. Sangiovanni-[56] Vincentelli, "Period optimization for hard real-time distributed automotive systems," in ACM DAC, 2007, pp. 278-283.
- K. Jiang, P. Eles, and Z. Peng, <sup>4</sup>Optimization of secure embedded systems with dynamic task sets," in *DATE*, 2013, pp. 1765–1770. [57]
- [58] S. Mohan, M.-K. Yoon, R. Pellizzoni, and R. B. Bobba, "Real-time systems security through scheduler constraints," in Euromicro ECRTS, 2014, pp. 129-140.
- [59] R. Pellizzoni, N. Paryab, M.-K. Yoon, S. Bak, S. Mohan, and R. B. Bobba, "A generalized model for preventing information leakage in hard real-time systems," in IEEE RTAS, 2015, pp. 271–282.
- [60] M.-K. Yoon, S. Mohan, C.-Y. Chen, and L. Sha, "TaskShuffler: A schedule randomization protocol for obfuscation against timing inference attacks in realtime systems," in *IEEE RTAS*, 2016, pp. 1–12. [61] K. Krüger, M. Völp, and G. Fohler, "Vulnerability analysis and mitigation
- of directed timing inference based attacks on time-triggered systems," in EUROMICRO ECRTS, vol. 106, 2018, pp. 22:1-22:17.
- [62] S. Mohan, S. Bak, E. Betti, H. Yun, L. Sha, and M. Caccamo, "S3A: Secure system simplex architecture for enhanced security and robustness of cyber-physical systems," in ACM international conference on High confidence networked systems. ACM, 2013, pp. 65-74.

- [63] D. Lo, M. Ismail, T. Chen, and G. E. Suh, "Slack-aware opportunistic monitoring for real-time systems," in *IEEE RTAS*, 2014, pp. 203–214.
  [64] F. Abdi, M. Hasan, S. Mohan, D. Agarwal, and M. Caccamo, "ReSecure: A restart-
- based security protocol for tightly actuated hard real-time systems," in IEEE CERTS, 2016, pp. 47–54.
  [65] F. Abdi, C.-Y. Chen, M. Hasan, S. Liu, S. Mohan, and M. Caccamo, "Guaranteed
- physical security with restart-based design for cyber-physical systems," in
- ACM/IEEE ICCPS, 2018, pp. 10–21.
  [66] R. Davis and A. Burns, "An investigation into server parameter selection for hierarchical fixed priority pre-emptive systems," in *IEEE RTNS*, 2008.
  [67] F. Abdi, J. 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 *IEEE CPSNA*, 2013, pp. 26–31.