

# 1 Mixed-Trust Computing for Real-Time Systems

2 **Dionisio de Niz**

3 SEI - Carnegie Mellon University, U.S.A.

4 dionisio@sei.cmu.edu

5 **Bjorn Andersson**

6 SEI - Carnegie Mellon University, U.S.A.

7 baandersson@sei.cmu.edu

8 **Mark Klein**

9 SEI - Carnegie Mellon University, U.S.A.

10 mk@sei.cmu.edu

11 **John P. Lehoczky**

12 Department of Statistics and Data Science - Carnegie Mellon University, U.S.A.

13 jl16@andrew.cmu.edu

14 **Amit Vasudevan**

15 SEI - Carnegie Mellon University, U.S.A.

16 avasudevan@sei.cmu.edu

17 **Hyoseung Kim**

18 Department of Electrical and Computer Engineering - U.C. Riverside, U.S.A.

19 hyoseung@ece.ucr.edu

20 **Gabriel A. Moreno**

21 SEI - Carnegie Mellon University, U.S.A.

22 gmoreno@sei.cmu.edu

## 23 — Abstract —

---

24 Verifying complex Cyber-Physical Systems (CPS) is increasingly important given the push to  
25 deploy safety-critical autonomous features. Unfortunately, traditional verification methods do not  
26 scale to the complexity of these systems. One promising approach to tackle this problem is a runtime  
27 verification variant [3] where small pieces of code (called *enforcers*) are added to the system to  
28 watch the system output replacing it if it is deemed incorrect. Verifying these enforcers leads to  
29 a system-wide correctness guarantee. A system built this way combines trusted components (the  
30 enforcers) with untrusted ones (the rest of the system). In order for this combination to work in a  
31 real-time system, it is necessary to (i) protect the trusted code from failures of the untrusted code  
32 both in terms of memory and time, and (ii) use a coordination mechanism between the trusted  
33 and untrusted parts in order to ensure that tasks meet their timing requirements. In this paper we  
34 present a new real-time scheduling framework to address these goals. This framework uses a new  
35 task model where each task is composed of an untrusted part guarded by a trusted one. We call  
36 them *mixed-trust tasks*. The untrusted part implements complex computations and has the support  
37 of a full OS running in a VM hosted by a trusted hypervisor. The trusted part, on the other hand,  
38 resides within the hypervisor. With this arrangement, if the untrusted part fails to finish (e.g., due  
39 to bugs) the trusted part is automatically activated to preserve safety properties (e.g., prevention of  
40 a crash). To support the new task model we present our new scheduling paradigm called *mixed-trust*  
41 *scheduling* that coordinates a fixed-priority non-preemptive scheduler protected by the hypervisor and  
42 a fixed-priority preemptive scheduler in the VM. We present a schedulability test and a procedure  
43 for configuring the parameters of the new task model. We also present our implementation based  
44 on the uberXMHF hypervisor running on a Raspberry Pi 3 and the Zero-Slack Rate-Monotonic  
45 (ZSRM) scheduler along with a coordination protocol designed to preserve trust. We present a set of  
46 experiments showing the system behavior during failures (transient and permanent) and a sample  
47 drone application to illustrate its use.

48 **2012 ACM Subject Classification** General and reference → General literature; General and reference



© Copyright 2019 Carnegie Mellon University, Hyoseung Kim and John P. Lehoczky. All Rights Reserved.;

licensed under Creative Commons License CC-BY



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

49 **Keywords and phrases** Dummy keyword

50 **Digital Object Identifier** 10.4230/LIPIcs...

## 51 **1 Introduction**

52 Certification authorities such as the FAA [23] allow the validation of different parts of a system  
53 with different degrees of rigor depending on their level of criticality. Formal verification has  
54 been recognized as a very important tool to provide correctness guarantees [2] for safety-  
55 critical components. Unfortunately, the verified property can be easily compromised if the  
56 verified components are not protected from unverified ones. Hence, **trust** requires that both  
57 verification and protection of components are jointly considered. This is the concept of trust  
58 we use in this paper.

59 One key limitation to building trust is the complexity of today’s operating systems that  
60 makes them impractical (if at all possible) to verify. As a result, there has been a trend  
61 to achieve a trusted computing base (TCB) by developing small verified hypervisors (HVs)  
62 and microkernels, e.g., seL4 [19], CertiKOS [16], and uberXMHF [28, 29]. In these systems,  
63 trusted and untrusted components co-exist on a single hardware platform but in a completely  
64 isolated and disjoint manner. We thus call it *disjoint-trust computing*. Trusted components  
65 in the TCB are typically made small and simple due to the difficulty of verification. They  
66 are strictly isolated from untrusted parts hosted in a virtual machine (VM) where rich  
67 functionalities are implemented on full-scale guest OSs like Linux.

68 The fundamental limitation of disjoint-trust computing is that it does not allow the  
69 use of untrusted components for critical functionality that rely on verification to assure  
70 safety. This is because the verified components must be isolated from the untrusted ones in  
71 a disjoint-trust computing platform if they are to be trusted. For instance, this prevents the  
72 use of untrusted machine learning algorithms<sup>1</sup> to drive a car if such a functionality needs  
73 to be verified. Instead, a separate trusted component would need to be in charge of the  
74 driving, isolating it from any untrusted component. Unfortunately, the complexity of the  
75 critical functionality demanded today, e.g., autonomous driving, makes the verification of  
76 these components very difficult or impossible.

77 In this paper, we propose *real-time mixed-trust computing* (RT-MTC) which enables the  
78 combination of trusted and untrusted components within critical functionality. Untrusted components  
79 are monitored by verified components, a.k.a. *enforcers*, ensuring that the output of the untrusted  
80 components always lead to safe states (e.g., avoid crashes). To ensure trust, enforcers are protected  
81 by a (verified) micro-hypervisor [29]. These enforcers are known as *logical enforcers* [3, 12]. In  
82 addition, RT-MTC includes *temporal enforcers*, which are small, self-contained codeblocks that perform  
83 default safety actions (e.g., *hover* in a quadrotor) if untrusted components have not produced a  
84 correct output by the deadline, thus preserving the temporal requirement of the safe behavior.  
85 Temporal enforcers are contained as extensions within the base (verified) micro-hypervisor without  
86 jeopardizing the existing level of trust (e.g., using compositional verification offered by extensible  
87 micro-hypervisors [29]).

88 Our framework incorporates two schedulers: (i) a preemptive fixed-priority scheduler in the VM  
89 to run the untrusted components and (ii) a non-preemptive fixed-priority scheduler within the HV<sup>2</sup>  
90 to run trusted components. To verify the timing correctness of safety-critical applications in our  
91 mixed-trust framework, we propose a new task model and schedulability analysis. We also present

---

<sup>1</sup> For which the research community has not yet devised any effective verification method.

<sup>2</sup> Recent advances in formal verification [ ] have shown that single-threaded verification is more tractable. This simplifies the verification of both the scheduler and the code the scheduler run within the HV.



92 the design and implementation of a coordination protocol between the two schedulers in order to  
 93 provide the protection and timing guarantees required by our framework. Lastly, we present an  
 94 implementation of our proposed framework using uberXMHF [29], an open-source, compositionally  
 95 verified micro-hypervisor framework supporting various security extensions. However, we note that  
 96 in principle, our framework can also be instantiated with other verified micro-kernels or hypervisors  
 97 provided they satisfy our requirements (see Section 2).

98 This work relies on innovations for code verification for the trusted parts that were presented in  
 99 previous publications. Since that is out of the scope of this paper, we refer the readers interested in  
 100 the compositional verification and isolation provided by uberXMHF to [29], the runtime verification  
 101 conducted by logical enforcers to [3], and the formal verification of temporal enforcement code to [6].

102 The remainder of this paper is organized as follows. Section 2 presents our RT-MTC framework,  
 103 an introduction to the runtime verification model that it supports, and the conditions that it must  
 104 fulfill to preserve the verified properties of the model. Section 3 gives the system model. Section 4  
 105 presents schedulability analysis of mixed-trust tasks, including the evaluation of the schedulability  
 106 analysis. Section 5 presents a fail-safe coordination protocol. Section 6 presents the implementation  
 107 of mixed-trust scheduling. Section 7 discusses the related work and Section 8 concludes.

## 108 **2 Mixed-Trust Real-Time Computing**

109 The *Mixed-Trust Real-Time Computing* framework is designed to provide both logical and timing  
 110 protection required to achieve trust in our previous runtime verification work [3]. In particular, a  
 111 system in [3] is modeled as a state machine with a set of states  $\mathbf{S}$  and a set of actions  $\Sigma$ ; when we  
 112 describe behavior, we let  $s$  in  $\mathbf{S}$  be a state and  $\alpha$  in  $\Sigma$ ; be an action. The evolution of the system is  
 113 modeled by the transition relation  $R$ , parameterized by the amount of time  $P$  that elapses during  
 114 the transition, and the action applied at the start of the transition. Formally  $R_P(\alpha) \subseteq \mathbf{S} \times \mathbf{S}$  is  
 115 the relation such that if the action  $\alpha$  is applied to the system at time  $t$  when it is in state  $s$  and  
 116 subsequently the system is in state  $s'$  at time is  $t + P$ , then  $(s, s') \in R_P(\alpha)$  — for a control-theoretic  
 117 analysis see [24]. We then define  $R_P(\alpha, s) = \{s' \mid (s, s') \in R_P(\alpha)\}$  as the set of states into which the  
 118 system can transition after taking action  $\alpha$ . We then identify  $\phi$  as the set of safe states. Given these  
 119 safe states we define a subset  $C_\phi$  of  $\phi$ -enforceable states as the largest set of states satisfying the  
 120 following two conditions:  $C_\phi \subseteq \phi$  and  $\forall s \in C_\phi \cdot \exists \alpha \in \Sigma \cdot R_P(\alpha, s) \subseteq C_\phi$ .

121 We denote by  $\text{SafeAct} : C_\phi \mapsto 2^\Sigma$  the mapping from  $\phi$ -enforceable states to actions that will  
 122 ensure that the system remains enforceable, i.e.,

$$123 \quad \text{SafeAct}(s) = \{\alpha \mid R_P(\alpha, s) \subseteq C_\phi\}$$

124 The action  $\alpha$  selected by the untrusted component in the system is then monitored and enforced  
 125 by the logical enforcer. The logical enforcer, defined as  $LE = (P, C_\phi, \mu)$ , receives  $\alpha$  from the  
 126 untrusted component. It is therefore assumed that the logical enforcer executes with the same period  
 127  $P$ .  $\mu(s) \in \text{SafeAct}(s)$  returns a set of enforcing actions. In each execution, the  $LE$  takes as input the  
 128 current system state  $s$  and the system action  $\alpha$  and produces an output action  $\tilde{\alpha}$  defined as follows:

$$129 \quad \tilde{\alpha} = \begin{cases} \alpha & \text{if } \alpha \in \mu(s) \\ \text{pick}(\mu(s)) & \text{otherwise} \end{cases} \quad (1)$$

130 where  $\text{pick}()$  selects one element from the set with whatever criteria. In the following we say  
 131 that  $\tilde{\alpha}$  is an  $LE$ -enforced action.

132 We now add to this model a temporal enforcer  $TE = (E, C_\phi, \alpha_T)$  that executes periodically  $E$   
 133 time units after the untrusted component job arrives, takes the enforced action  $\tilde{\alpha}$  from the  $LE$  and  
 134 generates a temporally-enforced action  $\hat{\alpha}$  before the end of the period as follows:

$$135 \quad \hat{\alpha} = \begin{cases} \alpha_T & \text{if } \tilde{\alpha} = \perp \\ \tilde{\alpha} & \text{otherwise} \end{cases} \quad (2)$$

136 where (i)  $\alpha_T \in \{\alpha \mid \alpha \in \text{SafeAct}(s) \forall s \in C_\phi\}$ , that is,  $\alpha_T$  is a safe action for any state in  $C_\phi$  (i.e.,  
 137 the specific state  $s$  is not needed to calculate  $\alpha_T$ ) and (ii)  $\perp$  denotes the absence of an action. In  
 138 the following we say that  $\hat{\alpha}$  is a  $TE$ -enforced action.



139 In this paper we present the software architecture, scheduling mechanisms, and schedulability  
 140 analysis to execute a taskset where each task includes untrusted code guarded by both a logical and  
 141 a temporal enforcer ensuring that not only the verification presented in [3] is preserved but also that  
 142 the system can be **trusted** to remain in the safe statespace region  $\phi$  according to our definition of  
 143 trust.

144 We now define the conditions that our framework must satisfy to support this model. For the  
 145 discussion of these conditions, we use *output* as the final action produced by the job once it has  
 146 been evaluated by the logical and temporal enforcers. These conditions are defined as follow:

- 147 ■ **C1.** Each task must produce an output every period.
- 148 ■ **C2.** There is only one output per period produced either by the *LE* or the *TE*.
- 149 ■ **C3.** A task should not produce both an output produced by the *TE* and another by the *LE*  
 150 within the same period.
- 151 ■ **C4.** An output produced by the task and validated by the *LE* must be the product of a  
 152 computation that executes within a single period, i.e., that reads the state of the system (e.g.,  
 153 senses), computes an output, and generates the output within the same period.
- 154 ■ **C5.** The *TE* of a task must execute  $E$  time units after the arrival of the job it guards and finish  
 155 before the end of the period.

156 To satisfy these conditions we not only need to create new runtime mechanisms, but the software  
 157 also needs to be structured in a way that takes advantage of these mechanisms. This is the topic of  
 158 our next section.

## 159 2.1 Software Architecture

160 Algorithm 1 shows the example behavior of a mixed-trust application. The **try** block shows the  
 161 core of the infinite loop that periodically senses the state, computes an actuation, and issues the  
 162 actuation. Within an iteration, the logical enforcer (*LE*) evaluates the computed actuation ( $\alpha$ )  
 163 and replaces it with a safe one ( $\tilde{\alpha}$ ) if needed (as presented in (1)). However, this loop can fail if  
 164 the code within the try block does not finish on time. Hence, a **catch** block is added to respond  
 165 to a timeout ( $E$  time units after the start of the current period). If the timeout occurs, then the  
 166 temporal enforcement actuation ( $\alpha_T$ ) is used to actuate. Note that it is not necessary to compute  
 167  $\alpha_T$  based on the current state given that it is safe in any state within  $C_\phi$ . Regardless of whether the  
 168 actuation done by the try or the catch, it is immediately followed by a wait for the completion of the  
 169 current period before executing another iteration. Note that in this algorithm there is no separate  
 170 temporal enforcer but it takes the form of the **try-catch** construct effectively implementing (2).

---

**Algorithm 1:** Behavior of a Mixed-Trust Periodic Task

---

```

1 while true do
2   try:
3      $s \leftarrow \text{currentState}()$ 
4      $\alpha \leftarrow \text{computeActuation}(s)$ 
5      $\tilde{\alpha} \leftarrow LE(s, \alpha)$ 
6     actuate( $\tilde{\alpha}$ )
7   catch timeout( $E$ ):
8     actuate( $\alpha_T$ )
9   end
10  waitForNextPeriod()
11 end

```

---

171 Now, even if the *LE* and the *TE* are formally verified, Algorithm 1 can still fail to preserve trust  
 172 in  $\phi$  if (i) the behavior of the *LE* is modified (once modified we do not consider that the output




© Copyright 2019 Carnegie Mellon University, Hyoseung Kim and John P. Lehoczky. All Rights Reserved.;

licensed under Creative Commons License CC-BY



Leibniz International Proceedings in Informatics

LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** Software Architecture

173 is from the *LE* — **C2**), (ii) the system fails to issue one of the actuations  $\tilde{\alpha}$  or  $\alpha_T$  before the end  
 174 of the period (**C1**), (iii) both  $\tilde{\alpha}$  and  $\alpha_T$  are issued within a period (**C3**), (iv) an  $\tilde{\alpha}$  calculated in a  
 175 previous period is issued (**C4**), or (v) the *TE* is modified (i.e., output is not considered to be a *TE*  
 176 output — **C2**).

177 Based on the runtime assurance requirements and the failure possibilities presented above, we  
 178 designed the software architecture presented in Figure 1. In this architecture, the green components  
 179 are trusted and need to be protected from untrusted components (in red). Note that the *LE* requires  
 180 the output of the controller ( $\alpha$ ) in order to calculate its output ( $\tilde{\alpha}$ ) as presented in (1). Hence,  
 181 while it can be (and must be) protected against logical behavior (code) modification, it cannot be  
 182 protected against delays given that the untrusted controller can choose to delay its output at will.  
 183 The *TE*, on the other hand, does not depend on  $\tilde{\alpha}$  since it only uses it to decide whether or not  
 184 to issue its safe action  $\alpha_T$ . However, the *TE* still needs to be protected against logical behavior  
 185 (code) modifications. Similarly, the communication of the  $\tilde{\alpha}$  from the *LE* to the *TE* must also be  
 186 protected against modification or falsification. Given this analysis, we define the following protection  
 187 requirements:

- 188 ■ **P1.** Logical behavior protection. This requires protecting both the code and the related internal  
 189 data. This is achieved through memory protection.
- 190 ■ **P2.** Temporal behavior protection. This requires protecting the activation time and the CPU  
 191 bandwidth allocated in order to meet real-time deadlines.
- 192 ■ **P3.** Communication authentication. This implies that we can verify who is the source and  
 193 ensuring that the source is protected (**P1**).
- 194 ■ **P4.** Communication logical integrity. This means that the message was not modified.
- 195 ■ **P5.** Communication temporal integrity. This means that the output generated is the product of  
 196 a computation within a period.

197 It is worth noting that the protections listed above are the protections of trusted code from  
 198 untrusted code from within the same task. Such requirements is a clear departure from other forms  
 199 of protection between different tasks. This led us to name this new task a **mixed-trust task** whose  
 200 timing characteristics will be formalized in Section 3. Before doing that, in the next section we  
 201 first discuss the design of our runtime architecture in order to satisfy these inter-task protection  
 202 requirements.

## 203 2.2 Runtime Architecture

204 As discussed in the introduction, we use a HV and its hosted VM to create a runtime environment  
 205 to execute different parts of mixed-trust tasks. In order to design this runtime environment, we first  
 206 design three protection domains to host the different parts, and a coordination protocol to preserve  
 207 the temporal behavior of the overall mixed-trust task.

208 The domains we designed are:

209 **Trusted Spatial protection Domain (TSD).** This is where the *LE* executes. It offers trusted  
 210 HV protection against memory modifications from the untrusted controller (or any other entity) but  
 211 does not offer temporal protection.

212 **Trusted Spatio-Temporal protection Domain (TSTD).** This is where the *TE* resides. It  
 213 offers trusted memory and temporal protection.



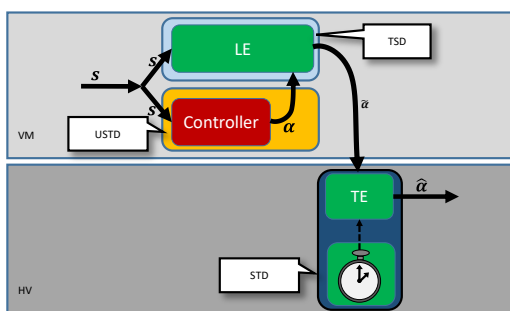
© Copyright 2019 Carnegie Mellon University, Hyoseung Kim and John P. Lehoczky. All Rights Reserved.;

licensed under Creative Commons License CC-BY



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 2** Runtime Architecture

214 **Untrusted Spatio-Temporal protection Domain (USTD)**. This is where the untrusted  
 215 controller resides. It offers untrusted VM spatial and temporal protection. This protection is  
 216 untrusted because it is implemented in the VM, which is assumed to not be verified. However, it is  
 217 expected its implementation is mature enough to work most of the time.

218 Figure 2 presents the location of these domains within the runtime architecture. This architecture  
 219 allows us to (i) minimize the code added into the HV space, (ii) protect the *LE* and hence the  
 220 integrity of the calculation of  $\tilde{\alpha}$  (**P4**), (iii) validate the  $\tilde{\alpha}$  origin by verifying the hypercall (syscall to  
 221 hypervisor) origin (**P3** – not shown in the figure), (iv) provide trusted logical protection for the  
 222 *TE* and the *LE* (**P1**), (v) provide trusted temporal protection for the *TE* (**P2**) and (vi) provide  
 223 untrusted temporal and spatial protection (**P1**, and **P2**) to the controller.

224 In order to guarantee **P5**, we added a scheduler in the HV and a coordination protocol that  
 225 synchronizes the scheduler in the VM with the one in the HV. Clearly, this coordination requires a  
 226 new integrated analysis that will be presented in Section 4. Hence, we defer the discussion of the  
 227 coordination protocol to Section 5. We now discuss the scheduling model.

### 228 3 System Model

229 Our system consists of a single-core processor<sup>3</sup> with a taskset  $\Gamma = \{\mu_i | \mu_i = (T_i, D_i, \tau_i, \kappa_i)\}$ . In the  
 230 taskset,  $\mu_i$  is a mixed-trust task with two execution segments,  $\tau_i$  and  $\kappa_i$ , with period  $T_i$  and deadline  
 231  $D_i$ . The segment  $\tau_i$  is considered to be untrusted and runs in the untrusted OS kernel inside the  
 232 VM. The segment  $\kappa_i$  is considered to be trusted code and runs within the trusted HV. To represent  
 233 the fact that these segments are handled by different schedulers, we consider them to be tasks and  
 234 call  $\tau_i$  the *guest task* (*GT*) and  $\kappa_i$  the *hyper task* (*HT*). These tasks are defined by:

$$235 \tau_i = (T_i, E_i, C_i) \tag{3}$$

236


$$237 \kappa_i = (T_i, D_i, \kappa C_i) \tag{4}$$

238 where  $T_i$  and  $D_i$  are the same as in  $\mu_i$ ,  $C_i$  is the worst-case execution time (WCET) of  $\tau_i$ , and  $\kappa C_i$   
 239 is the WCET of  $\kappa_i$ . Consider a particular job of  $\mu_i$ ,  $(\tau_{i,q}, \kappa_{i,q})$ . Ideally,  $\tau_{i,q}$  will execute correctly  
 240 taking no more than  $C_i$  time units and finishing within  $E_i$  time units of its activation. In this case,  
 241 the job  $\kappa_{i,q}$  is not activated. The logical enforcer (*LE*) verifies the correctness of  $\tau_{i,q}$ , while the  
 242 timing enforcer (*TE*) verifies the timing. If the logical enforcer (*LE*) does not notify the HV that  
 243  $\tau_{i,q}$  finished correctly and on time, then the corresponding HT  $\kappa_{i,q}$  is activated and runs at a higher  
 244 priority than any GT. The deadline for  $\tau_{i,q}$ ,  $E_i$ , is chosen to ensure that  $\kappa_{i,q}$  can finish by  $D_i$ , the  
 245  $\mu_i$  deadline. We show how to calculate  $E_i$  in Section 4.

246 Under our mixed-trust scheduling paradigm, HTs running in the HV have strictly higher criticality  
 247 than GTs in the VM. This leads to a scheduling hierarchy that strictly prioritizes all HTs over all

<sup>3</sup> For our current implementation we disable all cores but one in the Raspberry Pi 3. A multicore version is left for future work.





■ **Figure 3** Mixed-Trust Sample Timeline

248 GTs. That is, the execution of a HT is not preemptible by any GT running in the VM, and a  
 249 GT can be preempted by any HT that is ready to run. Furthermore, we assume that all HTs are  
 250 non-preemptive, whereas GTs can be preempted.

251 Under normal operation, the mixed-trust task  $\mu_i$  only runs its GT  $\tau_i$  executing at most  $C_i$  time  
 252 units, finishing within  $E_i$  time units and satisfying the  $LE$ , which then informs the scheduler of its  
 253 correct completion. However, if any of these three conditions fails (e.g., due to an error in  $\tau_i$  or  
 254 the untrusted OS kernel or a security infiltration), then its execution is interrupted and  $\kappa_i$  is run  
 255 within the HV. To detect this, a timer is set to expire  $E_i$  time units after  $\tau_i$ 's arrival. The goal of  
 256 the schedulability analysis is to compute the  $E_i$ s in order to ensure that all GTs can finish by  $D_i$   
 257 if all GTs execute correctly, and all activated HTs can finish by  $D_i$  if their corresponding GTs do  
 258 not complete correctly, that is the deadline of every task is met no matter whether any of the GTs  
 259 execute correctly or not. Fig. 3 depicts a sample execution timeline for a mixed-trust taskset with  
 260 three tasks. In this figure, HP/MP/LP represents high/medium/low priority on each scheduler, and  
 261 HC/LC denotes high/low criticality.

## 262 4 Schedulability Analysis

263 The schedulability analysis of a mixed-trust taskset is performed in three steps: we first calculate the  
 264 worst-case response time ( $R_i^k$ ) of each HT  $\kappa_i$  assuming non-preemptive fixed-priority scheduling<sup>4</sup>.  
 265 Then, we calculate  $E_i$  for each GT  $\tau_i$  by simply subtracting  $R_i^k$  from the deadline  $D_i$ . Finally, we  
 266 calculate the response time of each GT  $\tau_i$  and check whether it is at most  $E_i$ .

### 267 4.1 Hyper Task Response Time

268 To calculate the HT response time, we use previous results on non-preemptive fixed priority scheduling  
 269 (originally developed for the CAN bus) [11]. Specifically, the response time of a HT  $\kappa_i$  is calculated  
 270 in three steps:

- 271 1. Inspired by [11], we define the level- $i$  active period as a time interval in which the processor is  
 272 busy at all times and (i) there is at most one job from a task with lower priority than the priority  
 273 of task  $\kappa_i$  and if there is such a job, then it arrives before the beginning of the time interval, and  
 274 (ii) for the rest of the time interval, there is only execution of jobs from tasks with priority higher  
 275 than or equal to the priority of task  $\kappa_i$ . Our first step is to compute the maximum duration  
 276 of a level- $i$  active period. There are two reasons why we compute this maximum duration of  
 277 level- $i$  active period: (i) it allows us to compute the maximum number of jobs of task  $\kappa_i$  in the  
 278 level- $i$  active period, and (ii) we know that any execution outside the level- $i$  active period cannot  
 279 influence the response times of jobs of task  $\kappa_i$  in the level- $i$  active period.

<sup>4</sup> We assume unique priorities across the paper.



280 2. For a given HT  $\kappa_i$ , the start time of a job from  $\kappa_i$  in the level- $i$  active period is calculated. For  
 281 each job of  $\kappa_i$ , the finishing time of this job is then calculated by just adding the execution time.  
 282 This is because, once a task starts it cannot be preempted. Then, the response time of a job is  
 283 calculated as the finishing time minus its arrival time.

284 3. For a given HT  $\kappa_i$ , the response time is computed from the maximum response time across all  
 285 jobs of  $\kappa_i$  in the level- $i$  active period.

286 Let  $t_i^\kappa$  denote the maximum duration of a level- $i$  active period. Applying ideas in [11] on our  
 287 model yields that  $t_i^\kappa$  can be calculated using:

$$288 \quad t_i^\kappa = \max_{j \in \kappa L_i} \kappa C_j + \left\lceil \frac{t_i^\kappa}{T_i} \right\rceil \kappa C_i + \sum_{j \in \kappa H_i} \left\lceil \frac{t_i^\kappa}{T_j} \right\rceil \kappa C_j, \quad (5)$$

289 where  $\kappa L_i$  is the set of all HTs with lower priority than  $\kappa_i$  and  $\kappa H_i$  is the set of tasks with higher  
 290 priority than  $\kappa_i$ .

291 Given that a lower-priority task may be running when a higher-priority task arrives, (5) takes  
 292 into account the maximum interference from one job of a lower-priority task.

293 Let  $w_{i,q}^\kappa$  denote the latest starting time of the  $q^{th}$  job of  $\kappa_i$  in the level- $i$  active period. Then,  
 294 applying ideas in [11] on our model, we can obtain  $w_{i,q}^\kappa$  using:

$$295 \quad w_{i,q}^\kappa = \max_{j \in \kappa L_i} \kappa C_j + (q-1)\kappa C_i + \sum_{j \in \kappa H_i} \left( \left\lceil \frac{w_{i,q}^\kappa}{T_j} \right\rceil + 1 \right) \kappa C_j. \quad (6)$$

296 The response time can then be calculated as follows. For the jobs in the level- $i$  active period, we  
 297 can move the arrival times of the jobs to be as early as possible without violating the  $T$  parameters;  
 298 this may change the schedule but the duration of the level- $i$  active period is non-decreasing and the  
 299 starting time of each job is non-decreasing. Hence, it holds that the  $q^{th}$  job of each HT  $\kappa_j$  in the  
 300 active period (including  $\kappa_i$ ) arrives  $(q-1)T_j$  time units after the level- $i$  active period starts. For  
 301 each job of a HT  $\kappa_i$ , we can add  $\kappa C_i$  to its starting time and then subtract the arrival time of this  
 302 job and this yields the response time of the job. With these ideas, we can compute the response  
 303 time of  $\kappa_i$  as:

$$304 \quad R_i^\kappa = \max_{q \in \{1 \dots \left\lceil \frac{t_i^\kappa}{T_i} \right\rceil\}} (w_{i,q}^\kappa + \kappa C_i - (q-1)T_i). \quad (7)$$

305 It is worth noting that, in any schedulable taskset, the level- $i$  active period of a HT  $\kappa_i$  includes  
 306 only the execution of its first job if its corresponding GT  $\tau_i$  has a  $C_i > 0$ . This is because if the  
 307 taskset is schedulable, we verified that  $\tau_i$  has time to run for  $C_i$  and hence no HT runs at that time.  
 308 In other words, there is at least  $C_i$  time between two job executions (to completion) of a HT  $\kappa_i$   
 309 when  $\tau_i$  executes. Notwithstanding, we keep the equation that considers active tasks with multiple  
 310 job executions to allow tasksets without GTs.

311 Let us now discuss how to compute the  $E$  values. In order for HTs to be schedulable, clearly, we  
 312 must choose  $E$  values such that:

$$313 \quad \forall \mu_i \in \Gamma \quad E_i \leq D_i - R_i^\kappa. \quad (8)$$


314 We will later compute the response time of a GT  $\tau_i$  and denote it by  $R_i^g$ . In order for the schedulability  
 315 of a GT to hold, we must choose  $E$  values such that:

$$316 \quad \forall \mu_i \in \Gamma \quad R_i^g \leq E_i. \quad (9)$$

317 Ideally, we would like to develop an algorithm that computes  $E$ , for each task, such that if there  
 318 exists an assignment of values of  $E$  to tasks such that the schedulability test deems the taskset  
 319 schedulable, then our algorithm finds an assignment as well. In the appendix of the long version of  
 320 this paper [1], we present such an algorithm. Unfortunately, its time-complexity is high. Therefore,  
 321 we will calculate  $E_i$  in the following manner. For each  $\tau_i$ , we assign  $E_i$  to be as large as possible  
 322 while satisfying (8) and (9). The rationale for this is that choosing  $E_i$  to be as large as possible  
 323 gives as much space as possible for  $R_i^g$ ; however, it may adversely affect the  $R_i^g$  of a task  $\tau_l$  with  
 324 lower priority than  $\tau_i$ . Based on this discussion, we choose to calculate  $E_i$  using (10).

$$325 \quad E_i = D_i - R_i^\kappa. \quad (10)$$





■ **Figure 4** Aligning  $\tau_i$  and  $\tau_j$  leads to shorter  $R_{i,q}$  (case (a)) than aligning  $\kappa_i$  and  $\tau_j$  (case (b))

## 4.2 Guest Task Response Time

Recall that in the previous subsection, we defined the level- $i$  active period. Since we now consider response times of GTs, we need a new notion that is similar but also incorporates interference from HTs. Therefore, we define the level- $i$  busy period as a time interval such that at all times in the time interval, it holds that the processor is busy and there is only execution from jobs from HT or execution from jobs of GT of priority greater than or equal to the priority of  $\tau_i$ .

Next, we will prove how to determine the worst-case phasings. We will argue in two steps, first considering the case in which the task under analysis does not have any HT. The proofs are variants of Theorem 1 in [20].

► **Lemma 1.** *In the case when a mixed-trust task  $\mu_i$  does not have HT, the longest response time for all jobs of GT  $\tau_i$  occurs in a level- $i$  busy period initiated by the arrival of  $\tau_i$  and the arrival of other higher-priority GTs or HTs and lower-priority HTs.*

**Proof.** Following the argument of Lehoczky [20], let  $[0, b]$  denote a level- $i$  busy period in which  $\tau_i$  arrives at some point  $x_i$  after 0 during the busy period. Since the time before it arrives,  $[0, x_i)$ , is being used by higher-priority GTs, higher-priority HTs, or lower-priority HTs, moving its start to zero cannot change its completion time and can only increase its response time. Assume that a job of  $\tau_i$  arrives at 0. Also assume that some higher-priority GT,  $\tau_j$ , arrives after 0. From the definition of level- $i$  busy period, it follows that for all jobs of  $\tau_j$  that arrives strictly before 0, it holds that these jobs are not part of the level- $i$  busy period and also the corresponding HT jobs are not part of the level- $i$  busy period. Moving the arrival of  $\tau_j$  to 0 will result in an increased (or unchanged) amount of work in every interval  $[0, t]$  for every  $t$  in  $[0, b)$  possibly increasing or leaving unchanged the response time of  $\tau_i$  jobs. Similar arguments can be used if a higher- or lower-priority HT arrives after 0. ◀

Now consider the case in which the GT does have an associated HT. Aligning other task arrivals with the arrival of GT  $\tau_i$  does not necessarily cause the worst-case response for  $\tau_i$ —see Fig. 4.

► **Theorem 2.** *In the case when a mixed-trust task does have a HT, the longest response time for all jobs of GT  $\tau_i$  occurs in a level- $i$  busy period for which the following holds:*

1. *for each mixed-trust with higher priority than  $\mu_i$  it holds that either (i) the GT of the mixed-trust task arrives when the level- $i$  busy period starts or (ii) the HT of the mixed-trust task arrives when the level- $i$  busy period starts.*
2. *for the mixed trust task  $\mu_i$  it holds that either (i) the GT of the mixed-trust task arrives when the level- $i$  busy period starts or (ii) the HT of the mixed-trust task arrives when the level- $i$  busy period starts.*
3. *for each mixed-trust task with lower priority than  $\mu_i$ , it holds that the HT of the mixed-trust task arrives when the level- $i$  busy period starts.*

**Proof.** The sequence of jobs in the level- $i$  busy period will include a subsequence consisting of  $\tau_i$  and  $\kappa_i$  jobs. If the subsequence begins with a job of  $\tau_i$ , then Lemma 1 directly applies. Assume the subsequence begins with a job of  $\kappa_i$ .



© Copyright 2019 Carnegie Mellon University, Hyoseung Kim and John P. Lehoczky. All Rights Reserved.;

licensed under Creative Commons License CC-BY



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

364 Again let  $[0, b]$  denote a level- $i$  busy period. Assume that  $\kappa_i$  arrives at some point in time after  
 365 0 during the level- $i$  busy period. Unlike the previous case moving the start of  $\kappa_i$  to zero can reduce  
 366 its response time since it is non-preemptible, but it can still only increase the response times of all  
 367 of the GTs in the level- $i$  busy period.

368 Assume that  $\kappa_i$  arrives at 0, but some higher-priority GT,  $\tau_j$ , arrived after 0 but the previous  
 369 job of  $\kappa_j$  finished executing before 0 (hence it is not part of the busy period). Again moving the  
 370 arrival of  $\tau_j$  to 0 will result in an increased (or unchanged) amount of work in every interval  $[0, t]$  for  
 371 every  $t$  in  $[0, b)$  and therefore increase or leave unchanged the response time of all GT jobs in the  
 372 level- $i$  busy period. Similar arguments can be used if a higher or lower priority HT arrive after 0. ◀

373 This theorem identifies the phasings that need to be explored to determine a GT's worst-case  
 374 response time and motivate the equations that follow.

375 Previous work [5] has found the notion of a *request-bound function* useful in schedulability  
 376 analysis. The request-bound function for a set of jobs from a given task  $\tau_i$ , for a given time interval,  
 377 is the sum of the execution time of the jobs that have arrival times in that time interval. The  
 378 request-bound function for a given task  $\tau_i$ , for a given duration, is the maximum request-bound  
 379 function that jobs of this task can generate in a time interval of this duration. Given its usefulness,  
 380 we will now discuss how to create a request-bound function for our model; we will actually create a  
 381 notion of request-bound function that takes additional parameters. Also, our notion of request-bound  
 382 function gives the amount of execution of a mixed-trust task.

383 Recall that in our model, a task can generate a job but later the same job can “arrive” again ( $E$   
 384 time units later) to perform HV execution. Therefore, from the perspective of the request-bound  
 385 function, this arrival of HV execution is treated as the arrival of a job. The normal request-bound  
 386 function takes only two parameters: a task and a duration. In our model, we will use a more  
 387 specialized variant that takes two additional parameters,  $y$  (a phasing) and  $b$  (a 0-1 variable). We use  
 388 the former parameter ( $y \in \{E, A\}$ ) to indicate the phasing of the mixed-trust task  $\mu_i$ ; if  $y = E$ , then  
 389 we are computing the request-bound function for the phasing when the level- $i$  busy period starts at  
 390 a time when a HT of  $\mu_i$  arrives; analogously if  $y = A$ , then we are computing the request-bound  
 391 function for the phasing when the level- $i$  busy period starts at a time when a GT of  $\mu_i$  arrives. We  
 392 use the latter parameter ( $b \in \{0, 1\}$ ) to indicate whether the GT execution should be included in the  
 393 execution counted in the request-bound function.

394 The definition of request-bound function for our model is as given by the equation below:

$$395 \text{rbf}_i^{xy}(t, b) = \begin{cases} \left\lceil \frac{t - (T_i - E_i)}{T_i} \right\rceil^+ C_i b + \left\lceil \frac{t}{T_i} \right\rceil \kappa C_i & \text{if } y = E \\ \left\lceil \frac{t}{T_i} \right\rceil C_i b + \left\lceil \frac{t - E_i}{T_i} \right\rceil^+ \kappa C_i & \text{if } y = A \end{cases} \quad (11)$$

396 In the above definition, we use  $[x]^+$  to mean  $\max(0, x)$ .

397 We will use this notion of request-bound function to compute the response time of the GT  
 398 execution of a given mixed-trust task  $\mu_i$ . Then, if it holds that the computed response time for each  
 399 GT is less than or equal to its  $E$  parameter, then the taskset is schedulable (assuming that we have  
 400 already checked HT schedulability). Therefore, our goal is now to present equations for computing  
 401 the response time of a given GT. We will do so by presenting an equation for the maximum duration  
 402 of a level- $i$  busy period. Then, we compute the latest possible finishing time of a given job from a  
 403 given task in this level- $i$  busy period; then we also show that arrival times of jobs can be moved  
 404 to be as early as possible given the model; these two together (the finishing time and arrival time)  
 405 allow us to compute the response time of a GT job. Since we know the maximum duration of a  
 406 level- $i$  busy period, we can compute an upper bound on the number of jobs of a given task in a  
 407 level- $i$  busy period; we can compute the maximum response time over all these jobs of the given  
 408 task. This yields the GT response time. We will compute the GT response time for two cases: the  
 409 case that the given mixed-trust task arrives at the beginning of the level- $i$  busy period and the case  
 410 that the given mixed-trust task arrives with its HT aligned with the beginning of the level- $i$  busy  
 411 period. Given this high-level outline, we will now present the actual equations.

412 For each  $\mu_i$ , for each  $x \in \{E, A\}$ , let  $t_i^{g,x}$  denote the maximum level- $i$  busy period such that this  
 413 level- $i$  busy period starts with a job of the HT or the GT of  $\mu_i$  arriving ( $x$  indicates which). Then,



© Copyright 2019 Carnegie Mellon University, Hyoseung Kim and John P. Lehoczky. All Rights Reserved.;

licensed under Creative Commons License CC-BY



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

414 in a similar spirit as (5), we can, for  $x \in \{E, A\}$ , for a given task  $\tau_i$ , compute  $t_i^{g,x}$  as follows:

$$415 \quad t_i^{g,x} = \left( \sum_{j \in L_i} \text{rbf}_j^E(t_i^{g,x}, 0) \right) + \text{rbf}_i^x(t_i^{g,x}, 1) \\ + \sum_{j \in H_i} \max_{y \in \{E, A\}} \text{rbf}_j^y(t_i^{g,x}, 1). \quad (12)$$

416 Given a task  $\tau_i$  and a level- $i$  busy period, we refer to job  $q$  as the  $q^{\text{th}}$  job with a GT arrival in  
417 the level- $i$  busy period. For each  $\tau_i$ , for each  $x \in \{E, A\}$ , let  $w_{i,q}^{g,x}$  denote the maximum finishing  
418 time of job  $q$  of task  $\tau_i$ , relative to the start of the maximum level- $i$  busy period, such that this  
419 level- $i$  busy period starts with a job of the HT or the GT of  $\tau_i$  arriving ( $x$  indicates which). Then,  
420 in a similar spirit as (6), we can, for  $x \in \{E, A\}$ , for a given task  $\tau_i$ , for a given job index  $q$  of task  
421  $\tau_i$ , compute  $w_{i,q}^{g,x}$  as follows:

$$422 \quad w_{i,q}^{g,x} = \left( \sum_{j \in L_i} \text{rbf}_j^E(w_{i,q}^{g,x}, 0) \right) + qC_i + (q - 1 + I_{(x=E)})\kappa C_i \\ + \sum_{j \in H_i} \max_{y \in \{E, A\}} \text{rbf}_j^y(w_{i,q}^{g,x}, 1). \quad (13)$$

423 In (13),  $I_\phi$  is an indicator function that returns 1 if the Boolean predicate  $\phi$  is true and 0 otherwise.

424 For each  $\tau_i$ , for each  $x \in \{E, A\}$ , let  $R_{i,q}^{g,x}$  denote the maximum response time of job  $q$  of  $\tau_i$ ,  
425 relative to the start of the maximum level- $i$  busy period, such that this level- $i$  busy period starts  
426 with the arrival of a job of the HT or the GT of  $\tau_i$  ( $x$  indicates which). Then, in a similar spirit as  
427 part of (7), we can, for  $x \in \{E, A\}$ , for a given task  $\tau_i$ , for a given job index  $q$  of task  $\tau_i$ , compute  
428  $R_{i,q}^{g,x}$  as follows:

$$429 \quad R_{i,q}^{g,x} = w_{i,q}^{g,x} - ((q - 1)T_i + I_{(x=E)}(T_i - E_i)). \quad (14)$$

430 For each  $\tau_i$ , for each  $x \in \{E, A\}$ , let  $R_{i,q}^x$  denote the maximum response time of  $\tau_i$ , such that this  
431 level- $i$  busy period starts with the arrival of a job of HT or GT of  $\tau_i$  ( $x$  indicates which). Then, in a  
432 similar spirit as part of (7), we can, for  $x \in \{E, A\}$ , for a given task  $\tau_i$  compute  $R_i^{g,x}$  as follows:

$$433 \quad R_i^{g,x} = \max_{q \in \left\{ 1 \dots \left\lceil \frac{t_i^{g,x} - I_{x=E}(T_i - E_i)}{T_i} \right\rceil \right\}} R_{i,q}^{g,x} \quad (15)$$

434 Finally, the max response time of a GT over all phasings is obtained using:

$$435 \quad R_i^g = \max_{x \in \{E, A\}} R_i^{g,x} \quad (16)$$

### 436 4.3 Monotonicity and how to solve the equations

437 We must address the solution of these equations to determine schedulability. We consider three  
438 cases depending on utilization.

$$439 \quad 1. \sum_{\mu_i \in \Gamma} \frac{C_i + kC_i}{T_i} > 1$$

440 For this case, one can see that there is no solution to (12); the processor can be busy forever.

441 Hence,  $q$  can be arbitrarily large; for arbitrarily large  $q$ , (13) implies that  $w_{i,q}^{g,x}$  becomes infinite.

442 This does not necessarily imply that the taskset is unschedulable but it can also be seen that for  
443 this case ( $\sum_{\mu_i \in \Gamma} \frac{C_i + kC_i}{T_i} > 1$ ), for arbitrarily large  $q$ ,  $R_{i,q}^{g,x}$  becomes infinite as well. Thus, such

444 tasksets are unschedulable. Hence, if, for a taskset, it holds that  $\sum_{\mu_i \in \Gamma} \frac{C_i + kC_i}{T_i} > 1$ , then we  
445 terminate the schedulability analysis and report *unschedulable*.

$$446 \quad 2. \sum_{\mu_i \in \Gamma} \frac{C_i + kC_i}{T_i} < 1$$

447 For this case, one can see that (12) has a solution; hence, there is an upper bound on  $q$ . Then our

448 schedulability analysis will terminate. Note that many of our equations are of the form  $z=f(z)$ .

449 This is true for (5), (6), (12), and (13). For each of these equations, it holds that the expression on  
450 the right-hand side is monotonically non-decreasing in the variable on the left-hand side. Hence,  
451 these equations can be solved with fixed-point iteration.



© Copyright 2019 Carnegie Mellon University, Hyoseung Kim and John P. Lehoczky. All Rights Reserved.;

licensed under Creative Commons License CC-BY



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

452 3.  $\sum_{\mu_i \in \Gamma} \frac{C_i + \kappa C_i}{T_i} = 1$

453 For this case, it is difficult to determine whether (12) has a solution. We will pessimistically  
 454 report *unschedulable* for such tasksets.

#### 455 4.4 Generalizations

456 It is worth noting that a taskset is not required to have tasks with both GT and HT. More  
 457 importantly, when a taskset contains no HT our scheduling equations reduce to fixed-priority  
 458 preemptive response time analysis. Similarly, when a taskset contains no GT, then it reduces to  
 459 fixed-priority non-preemptive schedulability analysis.

460 We can also observe these properties realized in a running system. For example, even if GTs  
 461 have their corresponding HTs, the system will run like a “preemptive” scheduling system if all GTs  
 462 finish without exceeding their  $C_i$ . On the other hand, if a VM crashes, the system will run like a  
 463 pure non-preemptive scheduling system as will be shown in Figure 10 of Section 6.4.

#### 464 4.5 Enforcement

465 Given that GTs are not trusted, their  $C_i$  values need to be enforced. This enforcement allows us  
 466 to implement a graceful degradation scheme by preventing failing GTs from interfering with other  
 467 non-failing GTs. Clearly, if a failure affects the kernel in the VM (e.g., due to a security attack) all  
 468 the GTs will be compromised but the HV and the HT will be protected from the failure. In contrast  
 469 to the GTs, the HTs are trusted and their  $\kappa C_i$  does not need to be enforced. In addition, there can  
 470 be two possible GT enforcement options when the enforcement timer elapses: (i) the execution of  
 471 the job of a GT  $\tau_i$  is aborted immediately and the corresponding HT  $\kappa_i$  is responsible for cleaning  
 472 up its execution, or (ii) the job of the GT  $\tau_i$  is deferred (suspended) and its HT  $\kappa_i$  executes only  
 473 temporary actions (e.g., safe actuation in a control task), allowing the GT’s job to complete in the  
 474 next period. Our current implementation uses the latter option and we will call it *deferral GT*  
 475 *enforcement*.

#### 476 4.6 Experiments

477 This section presents experiments that show how taskset parameters influence schedulability. We  
 478 found three reasons that impact schedulability:

- 479 1. When considering one task, its HT can experience interference from HT of other tasks. In addition,  
 480 the HT of other tasks can also delay the execution of the GT of the task under consideration.  
 481 This is a double-accounting effect that has no analog in classic fixed-priority scheduling.
- 482 2. A GT can experience a long delay because of execution of HT of all the other tasks. This reason  
 483 is most impactful when the GT has very small period.
- 484 3. Consider the lowest-priority task and consider its GT. When its period is around  $\sqrt{2}$  of the  
 485 period of its higher priority task, then the schedulability deteriorates (just like it does for classic  
 486 rate-monotonic preemptive scheduling).

487 We now illustrate the schedulability conditions we just introduced with the following experiments.  
 488 In these experiments we vary: (1) the taskset utilization, (2) the ratio between the maximum period  
 489 and the minimum period, (3) the number of tasks in the taskset, (4) the ratio between the HT  
 490 WCET and the sum of the GT and HT WCET, and (5) the ratio between the deadline and the  
 491 period of the tasks. The ranges of variations are presented in Table 1.


492 We perform five experiments to vary utilization,  $\frac{T_{max}}{T_{min}}$  ratio, number of tasks,  $\frac{\kappa C}{C + \kappa C}$  ratio, and  
 493  $\frac{D}{T}$  ratio.

494 The default values for the parameters that do not vary are presented in Table 1. Two observations  
 495 are in order. First, the default number of tasks is set to 10 given that a larger number of tasks  
 496 reduces the chance of having a schedulable taskset as can be seen in Fig. 6a. Second, the default  
 497 utilization is set to 80% also to reduce the influence of the utilization to dominate when varying the  
 498 other parameters.



Parameter	Range	Default
Number of Tasks	{3, 4, ..., 200}	10
$U$	{0.1, 0.2, ..., 1.0}	0.8
$\frac{\kappa C}{C + \kappa C}$	{0.1, 0.2, ..., 1.0}	0.1
$\frac{T_{max}}{T_{min}}$	{1, 2, 4, ..., 1024}	100.0
$\frac{D}{T}$	{0.1, 0.2, ..., 1.0}	1.0
$T_{min}$		1000

■ **Table 1** Parameter Ranges and Defaults



(a) As utilization grows

(b) As  $\frac{T_{max}}{T_{min}}$  grows


■ **Figure 5** Success rate as utilization and  $\frac{T_{max}}{T_{min}}$  grow

499 In the experiments each data point is the average of 100,000. Specifically, each dot is computed  
500 from 100,000 tasksets and we compute the percentage of schedulable tasksets.

501 Each taskset is generated with the selected number of tasks. Each task is assigned a utilization  
502 equal to the selected total utilization of the taskset divided by the number of tasks. Then the period  
503 of the task is chosen at random (with uniform distribution) from the period range selected.

504 Fig. 5a shows the percentage of schedulable tasksets as the taskset utilization grows from 10%  
505 to 100%. Note that the experiment shows a decline in the percentage of schedulable taskset just  
506 after 20%. This is due to reason 2.

507 Fig. 5b depicts the percentage of schedulable tasksets as the ratio of the minimum and maximum  
508 period grows. Here, we can see that when increasing the ratio, the success rate decreases and then  
509 increases and then decreases again. The initial decrease is caused by reason 3; the second decrease  
510 is caused by the reason 2.



(a) As number of tasks grow

(b) As  $\frac{\kappa C}{C + \kappa C}$  grows

■ **Figure 6** Success rate as as number of tasks and  $\frac{\kappa C}{C + \kappa C}$  grow




© Copyright 2019 Carnegie Mellon University, Hyoseung Kim and John P. Lehoczky. All Rights Reserved.;

licensed under Creative Commons License CC-BY



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 7** Success rate as  $\frac{D}{T}$  grows

511 Fig. 6a presents the fraction of schedulable tasksets as the number of tasks in the taskset grows.  
 512 The curve decreases exponentially reaching zero at about 115 tasks. This is because when the  
 513 number of tasks increase, the maximum period divided by the minimum period among the tasks  
 514 generated becomes larger and then reason 2 becomes more impactful.

515 Fig. 6b shows schedulable tasksets percentage as the ratio of HT WCET to the combined HT  
 516 and GT WCET grows. The figure shows a quick drop in the percentage of schedulable tasksets as  
 517 this ratio increases. This is because of reason 2.

518 Fig. 7 presents the percentage of schedulable tasksets as the ratio of deadline to period grows.  
 519 As expected, when deadlines are much smaller than periods tasks are not schedulable. The situation  
 520 starts to improve after a ratio of 0.6 when we start finding schedulable tasksets. Then the percentage  
 521 of schedulable tasksets grows fast, and the rate of improvement reduces at the end of the ratio range.

## 522 **5 Fail-Safe Mixed-Trust Scheduling Coordination Protocol**

523 With the timing analysis as a background we can now discuss the coordination protocol between our  
 524 schedulers. A key challenge that needed to be solved in our framework was the prevention of any  
 525 dependency of trusted HV and HT code from the untrusted code, while still enabling the successful  
 526 coordination of the HV and guest schedulers. The dependency of a higher-critical component from a  
 527 lower-critical one is known as dependency inversion [14].

528 Preventing the dependency inversion problem necessitates three mechanisms: (1) a Secure HT  
 529 Bootstrapping (*SHTBoot*) Protocol, (2) a Fail-Safe HT Triggering (*FSHTTrigger*) mechanism, and (3)  
 530 a Late-Output Prevention protocol (*LOP*).

### 531 **5.1 Secure HT Bootstrapping (*SHTBoot*)**

532 The objective of the SHTBoot protocol is to ensure that the HT can start and execute periodically  
 533 according to its specification even if the VM is unable to bootstrap the GT. This is in fact necessary  
 534 to properly implement the trusted temporal protection of the **TSTD** and the protection requirement  
 535 **P2** (discussed in Section 2). We leverage secure boot mechanism provided by uberXMHF to ensure  
 536 that the micro-hypervisor framework is the first to get control when the system is powered on.  
 537 The SHTBoot protocol starts the HTs and GTs independently out of bootstrapping task tables  
 538 stored in the HT and VM storage, respectively. To synchronize the periodic activation of a GT with  
 539 the corresponding enforcement timer of its HT, the guest scheduler requests the start time of the  
 540 next period from the HV and uses this time to start the first job of the GT, aligning the periodic  
 541 activations of the GTs and HTs as required.

### 542 **5.2 Fail-Safe HT Triggering Protocol (*FSHTTrigger*)**

543 The objective of the FSHTTrigger is to prevent a failure in the VM from disabling or corrupting the  
 544 periodic activation of the HTs. This is in fact the activation side of the protection requirement **P2**.



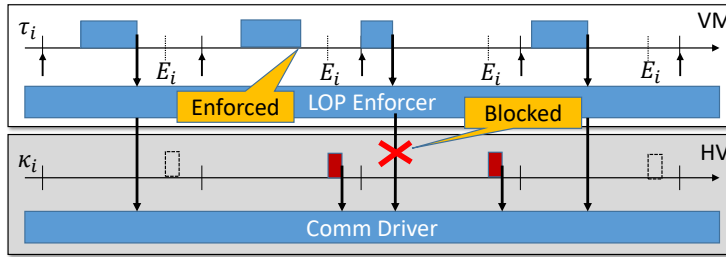
© Copyright 2019 Carnegie Mellon University, Hyoseung Kim and John P. Lehoczky. All Rights Reserved.;

licensed under Creative Commons License CC-BY



Leibniz International Proceedings in Informatics

LIPIC Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 8** Late-Output Prevention Protocol

545 To implement the FSHTrigger, a general strategy followed is to program separate timers for the  
 546 VM and the HV down to the hardware level so that untrusted VM code can program its own timer  
 547 but cannot program the HV timer. This way when a task  $\mu_i$  is started, its  $E_i$  timer is programmed  
 548 within the HV and will always trigger no matter what code executes in the VM (even malicious  
 549 code actively trying to disable the timer). As a result, the HT can always run on time and complete  
 550 its safety action by the deadline. We leverage the peripheral isolation provided by uberXMHF to  
 551 isolate the HV and guest timers so that the guest cannot access the HV timer.

552 Even though an HT must be isolated from failures of the corresponding GT, some of their  
 553 timing parameters need to be synchronized: (1) initial release offset, which is needed to program an  
 554 enforcement timer exactly  $E_i$  units after the arrival of each job, and (2) job completion time, which  
 555 is to disable the HT if it completed before  $E_i$ . The creation time synchronization is performed by  
 556 SHTBoot. For the completion-time synchronization we rely on the logical enforcer in the VM to send  
 557 the completion signal to the HV. More specifically, the logical enforcer ( $LE$ ) verifies that the output  
 558 produced is safe (or replaces it with a safe one) and sends the completion signal. To guarantee  
 559 that the  $LE$  only sends the completion signal when a safe output is generated, the HV protects the  
 560  $LE$  memory and we assume that the  $LE$  code's trustworthiness has been asserted (e.g. through  
 561 verification). This satisfies the protection requirement **P1** and **P3**. The only possible failure is then  
 562 a denial-of-service, i.e., **P2** is not guaranteed. In particular, if the GT  $\tau_i$  takes longer than  $C_i$  to  
 563 complete, the task will be enforced and the  $LE$  will neither complete nor send the completion signal  
 564 to the HV before  $E_i$ . However, this failure is part of the assumption of the FSHTrigger protocol  
 565 since the absence of the completion signal will trigger the HT and issue the safe output. In other  
 566 words, the HT that host the  $TE$  preserves its temporal behavior (**P2**).

### 567 5.3 Late-Output Prevention ( $LOP$ ) Protocol

568 A late output may occur when a GT job is allowed to complete and generates an output (e.g.,  
 569 actuation commands)  $E_i$  time units or more after its arrival. Recall the deferral GT enforcement  
 570 approach explained in Section 4.5. In this case, a job can be suspended once  $E_i$  time units have  
 571 elapsed after its arrival, and resumed in the next period, allowing it to send its output in the second  
 572 period (violating requirement **C4**). Preventing this output is important because the logic in the  
 573 application algorithm (e.g. control algorithm) assumes that it is computed within the execution of  
 574 a single job, perhaps using inputs (sensing) from the beginning of the period that are only valid  
 575 for output (e.g. actuation) during this same period (**C4**). To solve this problem, the output and  
 576 completion signals are bundled in a single call and all the output is mediated by an LOP enforcer in  
 577 the VM kernel scheduler, making sure that the output is discarded if it is sent after  $E_i$ . We also  
 578 protect the LOP enforcer memory (in protection domain **TSD**) in the kernel and assume that the  
 579 trust in its code has been established (e.g. through verification). As a result, the LOP enforcer can  
 580 only fail by not sending the output. The use of the LOP enforcer separately from the  $LE$  allows  
 581 us to separate the logical from the temporal correctness, simplifying its verification but preserving  
 582 both properties.

583 Figure 8 depicts an example of the execution timeline of the LOP protocol, where the second job  
 584 of GT  $\tau_i$  is budget-enforced and does not send the output and completion signal in that period. The  
 585 absence of this signal triggers the HT (red rectangle), which in turn sends the safe output. When  
 586 the GT is resumed in the next period, it produces an output but the LOP blocks it along with the



© Copyright 2019 Carnegie Mellon University, Hyoseung Kim and John P. Lehoczky. All Rights Reserved.;

licensed under Creative Commons License CC-BY



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

587 completion signal. Hence, the HT is activated again to send the safe output.

## 588 **6 Implementation**

589 Our scheduling mechanisms are implemented by the combination of the budget enforcement of  
590 the ZSRM kernel scheduler [13] running in a VM and a non-preemptive fixed-priority scheduler  
591 implemented within the verifiable and extensible uberXMHF micro-HV [28, 29, 27] running in a  
592 Raspberry Pi-3 platform. In order to prevent failures in the VM from affecting the code within the  
593 micro-HV, uberXMHF implements two-stage hardware memory page-tables and protections that  
594 cannot be modified by the guest OS running inside the VM. The two-stage hardware page tables  
595 isolate the micro-HV memory where the HTs reside. The guest OS memory isolation and protections  
596 have been formally verified as presented in [28, 29]. uberXMHF employs a compositional verification  
597 methodology that allows the addition of security sensitive functionality as modularly protected and  
598 verifiable components (called *uberapps*). We used this facility to implement the HV scheduler as a  
599 *uberapp*.

### 600 **6.1 Hyper-Task-Aware Budget Enforcement**

601 ZSRM is a Linux kernel module scheduler that works on top of the fixed-priority preemptive scheduler  
602 of the Linux kernel. The original ZSRM budget enforcement is performed by shadowing the priority  
603 queues of the fixed-priority scheduler to keep track of the current top-priority active task and the  
604 amount of CPU time this task consumes. Tasks become ready to run when they are created, and  
605 they go to sleep by calling the ZSRM API `wait_next_period()` when they finish their periodic job  
606 execution. At this time, the shadow priority queues in ZSRM are updated, and the top-priority  
607 active task in the queue is scheduled and marked as the `current` task. Similarly, when a task period  
608 elapses, ZSRM wakes the task up and makes a scheduling decision. Whenever a task is scheduled,  
609 an enforcement timer is set to expire with the maximum remaining budget of that task. If the task  
610 either finishes or is preempted by another task before the enforcement timer expires, the timer is  
611 canceled and reprogrammed for the next task. On the other hand, if the timer expires, the task is  
612 suspended until its next periodic timer expires. The budget accounting is implemented by recording  
613 a starting timestamp when a task becomes the `current` task and a finishing timestamp when the task  
614 is either preempted or completes its periodic execution. Then, subtracting the starting timestamp  
615 from the finishing one, gives us the CPU time used by the task. We accumulate this time for all the  
616 intervals that a task is considered to be the currently executing task in each period. Then the budget  
617 minus this accumulated CPU time consumption is used to program the enforcement timer every  
618 time the task becomes the currently executing task. The implementation of this budget enforcement  
619 mechanism was formally verified in [6].

620 Given that HT preemptions are invisible to the VM and ZSRM scheduler, the ZSRM budget  
621 enforcement fails to account for them. To handle this, we use an event logger within the HV to  
622 record the timestamps of the activation and completion of the HTs. Then, these events are requested  
623 from the HV when the budget enforcement timer triggers. At this time, we use the HT timestamps  
624 to discount the HT preemption time from the budget. Finally, if this discount reduces the budget  
625 consumption to less than its  $C_i$  budget, the budget enforcement timer is reprogrammed to allow the  
626 GT to keep executing for the difference.

### 627 **6.2 Implementation Overhead**

628 Table 2 shows the average overhead of our implementation in the Raspberry Pi-3 running at 1200  
629 MHz. A few observations are in order: (i) The context switch includes the shadow queue manipulation  
630 in the ZSRM kernel module but not the kernel scheduling queues. (ii) The budget enforcement  
631 includes the activation of a kernel thread that is used to prevent modifying the priority queues of  
632 the kernel from inside the timer interrupt context, as well as the hypercall to obtain the hypertask  
633 preemptions. The hypercall to get the hypertasks preemptions is also shown separately.



© Copyright 2019 Carnegie Mellon University, Hyoseung Kim and John P. Lehoczky. All Rights Reserved.;

licensed under Creative Commons License CC-BY




Leibniz International Proceedings in Informatics

LIPIC Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Procedure	Avg Overhead (ns)	WC Overhead (ns)
Job Arrival with context Switch	12368	14105
Job Arrival, no context switch	7789	9052
Budget Enforcement	26263	42894
wait_next_period()	7789	9894
Get Hypertask Preemptions	7473	11315

■ **Table 2** Overhead



■ **Figure 9** Mixed-Trust Taskset Execution Timeline (timestamps in  $10^{-4}$  seconds)

### 6.3 Spurious Temporal Failure Illustration

We ran a taskset of two tasks with our implementation to illustrate the mixed-trust scheduler. This taskset allowed us to obtain a plot of the tasks with timestamps captured from both the kernel scheduler and the HV scheduler. Both schedulers query the same hardware timer to read their timestamps, allowing us to have an integrated timeline without incurring context-switch penalties.

Figure 9 shows the execution of two mixed-trust tasks in a timeline reconstructed from timestamps for the following events: (i) **arrival (A)** that marks when the job of the task becomes ready to execute; (ii) **job finishing (F)** by calling `wait_next_period()` from a guest application task to the kernel scheduler when its job ends normally; (iii) **budget enforcement (B)**; (iv) **resume** marks the start of a colored rectangle showing when the job starts to execute; and (v) **paused** presented as the end of a colored rectangle. It is worth noting that the activation of the HTs is presented as a small rectangle marked by the **resume** and **paused** events and filled with wavy and checkered patterns. The timeline shows different types of preemptions as follows. From 1806 to 1888,  $\mu_1$ 's HT ( $\kappa_1$ ) preempts  $\mu_2$ 's GT ( $\tau_2$ ). Then from 2005 to 2205,  $\tau_2$  is preempted again this time by  $\tau_1$ . At 2205,  $\tau_1$  is budget-enforced letting  $\tau_2$  to run until it is enforced at 2380. Given that  $\tau_1$  did not signal completion, its HT  $\kappa_1$  executes from 2806 to 2888. Then the next job of  $\tau_1$  arrives at 3005 and executes until it is preempted by the HT  $\kappa_2$  at 3201.  $\kappa_2$  finishes its execution at time 3283 allowing  $\tau_1$  to resume execution until it is enforced at 3285. The last job execution shown in the timeline is a normal that start at 3400 and finishes at 3500.

### 6.4 Permanent Failure Illustration

In this section we present an experiment to show how our approach handles a complete failure of the kernel in the VM. More specifically, we start two tasks ( $\mu_1$  and  $\mu_2$ ) in this case running with a periodicity of  $T_1 = 1$  and  $T_2 = 2$  seconds, respectively, with their respective HTs designed to run for only 10 ms each. We let the task  $\mu_1$  run for four periods and  $\mu_2$  for three periods without faults (not even timing faults so their HT do not trigger). The third job of  $\tau_2$  uses a semaphore to signal a third task (not shown) whose only role is to wait for this signal and invoke a system call in our scheduler specifically designed to test a full kernel failure (this scheme allows  $\tau_2$  to properly finished given that the third task has the lowest priority). This call, in turn, calls the kernel `panic()` function. The `panic()` function is designed to completely stop the kernel when a part of the kernel reaches a point where it is not possible to continue executing due to unrecoverable failures. This basically simulates a hard crash. In order to capture the timestamps, we send them to a serial port through the HV. However, because the serial port driver implementation is slow, it creates some disruption in the timestamps. The resulting trace is presented in Figure 10.

A few observations about the trace in Figure 10 are in order. First, the call to `panic()` occurs




© Copyright 2019 Carnegie Mellon University, Hyoseung Kim and John P. Lehoczky. All Rights Reserved.;

licensed under Creative Commons License CC-BY

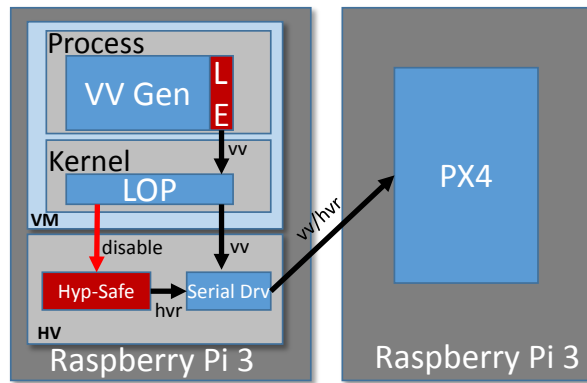


Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 10** Experiment with Permanent VM crash (timestamps in  $10^{-2}$  seconds)



■ **Figure 11** Drone Application Architecture

668 after  $\tau_2$  finishes at time 408. Second, after this time no more arrival (A) or finishing (F) events occur  
 669 from either of the tasks. Third, as expected both HTs ( $\kappa_1$  and  $\kappa_2$ ) continue executing periodically.  
 670 Finally, the first execution of both HTs after the `panic()` call occurs almost two periods from the  
 671 last GT executions. In particular, the first execution of  $\kappa_1$  after the `panic()` call occurs at 515 that  
 672 is almost two periods from the last arrival of its GT at 317. Similarly, for  $\kappa_2$  its corresponding first  
 673 arrival after the `panic()` call is at 806 and its last GT arrival at 408. This is expected because  
 674 the HT execution is scheduled at the end of the period of a task. This means that when the GT  
 675 executes at the beginning of the period, the execution of the HT of the following period will happen  
 676 almost two periods apart, even though an output is produced in every period.

## 677 6.5 Illustrative Application

678 We implemented a sample mixed-trust application of a drone mission. This application consists of  
 679 two components: (i) the mission controller, which generates velocity vectors (VV) that the drone  
 680 must fly to follow a route, and (ii) the Pixhawk<sup>5</sup> flight controller running the PX4 autopilot<sup>6</sup> in  
 681 off-board mode, which makes the drone fly in the direction and speed of the last VV received. The  
 682 mission controller runs in its own processor sending a VV message every 50 ms to the processor  
 683 where the flight controller runs. A logical enforcer in the mission controller prevents the drone from  
 684 violating spatial constraints (e.g., a virtual fence or the collision volume of other drones [7]). In  
 685 addition, we added an HT to the mission controller can take a safe action and continue to send VV  
 686 messages to the flight controller in case the GT fails.

687 Figure 11 depicts this structure, which, for simplicity of presentation, shows only one mixed-trust

<sup>5</sup> <https://pixhawk.org/>

<sup>6</sup> <http://px4.io/>



688 task. This task has a guest task that generates velocity vectors (*VV Gen*) and a hypertask (*Hyp-Safe*)  
689 that generates the safe drone action *hover* (*hvr*), which is a null VV. The figure also shows the LOP  
690 mediation of the messages sent by *VV Gen* ensuring that (a) no late outputs are allowed, and (b)  
691 when no output is generated by *VV Gen* the *Hyp-Safe* HT generates the *hover* action. The mission  
692 controller was implemented using a version of DronecodeSDK<sup>7</sup> that we modified to handle serial  
693 communications through the serial driver in the HV, and to use the bundled output and completion  
694 signals for the LOP. We ran this application using hardware-in-the-loop simulation (i.e., actual  
695 mission and flight computers connected to a drone simulation), which allowed us to observe the  
696 physical consequences. We tested both temporary failures and hard failures where we verified that  
697 both the LOP and the HTs properly prevent drone failures.

## 698 **7 Related work**

699 Previous work has recognized that small operating systems kernels can be more reliable [18] and can  
700 be formally verified [19]. In this context, decomposing an application can provide security benefits as  
701 well [26]. But they do not provide schedulability analysis. Previous work on hierarchical scheduling  
702 (e.g., [15]) studies run-time systems with two schedulers and they present theories that provide  
703 real-time guarantees but they do not consider a task that spans different components. Operating  
704 systems works considering real-time requirements have also been presented [21, 9, 4, 17, 30] to  
705 achieve isolation and some offer offline schedulability tests but not for the task model that we  
706 consider (where a task can span two operating systems) and they do not target formal verification  
707 of operating system code.

708 Previous work [22, 8] combining real-time and security are not based on a runtime verification  
709 framework that requires the integration of trusted and untrusted components.

710 Works on mixed-criticality (see [10] for an excellent survey) share our goal of monitoring run-time  
711 behavior and taking action when behaviors that are abnormal are detected; however, we are not  
712 aware of any work on a mixed-criticality scheduler that considers our task model and uses a formally  
713 verified kernel.

714 Simplex [25] is an architecture for designing controllers. It comprises one complex controller,  
715 a simple controller, and two sets of states. The first set describes the safe states; the second set  
716 describes when there is a need to transition between controllers. The complex controller is allowed  
717 to operate when the plant is in the second set. If the plant leaves this set, then the simple controller  
718 takes over. With this architecture, the complex controller can be optimized for performance and  
719 does not need to be verified; the simple controller, however, is verified to make sure that the plant is  
720 always in a safe state. One can think of the simple controller in Simplex as somewhat analogous to  
721 our HT.

## 722 **8 Conclusions**

723 In this paper we presented a real-time mixed-trust framework for mixed-trust tasks composed of  
724 an untrusted part (a.k.a. guest task (GT)) that runs in a VM guarded by a temporal enforcer  
725 (a.k.a. the hyper task (HT)) running in the trusted HV that hosts the VM. With this arrangement  
726 if the untrusted GT fails to complete on time, the HT is automatically triggered executing a safe  
727 computation to keep the system safe (e.g., avoiding a crash). We presented the new real-time  
728 schedulability analysis and parameter configuration algorithms along with experiments that explore  
729 the sensitivity of our analysis to different parameters. We also presented the implementation of  
730 our framework based on the coordination of a non-preemptive fixed-priority scheduler within the  
731 uberXMHF HV and the ZSRM scheduler in the VM. We also discussed the three coordination  
732 mechanisms used to preserve trust even in the presence of failures and presented experiments to  
733 show the behavior of our system during transient and permanent failures. Finally, we presented  
734 a sample drone application to illustrate the use of our framework. The contributions presented in

---

<sup>7</sup> <https://www.dronecode.org/>



735 this paper enable the use of complex non-deterministic software (e.g., machine learning) in safety-  
736 critical systems (including dependencies on large OS and libraries) while still preserving provable  
737 logical and timing guarantees. Furthermore, informed by runtime verification, our framework avoids  
738 taking actions that, while preserving timing properties, break logical properties, e.g., killing a job  
739 disregarding the consequences of the absence of actuation. This is, in fact, the result of the co-design  
740 of a logical and timing runtime verification scheme that gave birth to our mixed-trust scheduling  
741 framework.

## 742 **Acknowledgment**

743 Copyright 2019 Carnegie Mellon University, Hyoseung Kim and John P. Lehoczky. All Rights  
744 Reserved. This material is based upon work funded and supported by the Department of De-  
745 fense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation  
746 of the Software Engineering Institute, a federally funded research and development center. NO  
747 WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING  
748 INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNI-  
749 VERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED,  
750 AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR  
751 PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE  
752 OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WAR-  
753 RANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK,  
754 OR COPYRIGHT INFRINGEMENT.

755 [DISTRIBUTION STATEMENT A] This material has been approved for public release and unlimited  
756 distribution.

757 Please see Copyright notice for non-US Government use and distribution. Carnegie Mellon® is  
758 registered in the U.S. Patent and Trademark Office by Carnegie Mellon University.

759 DM19-0109

## 760 **References**

- 761 1 Mixed-Trust Scheduling. [https://www.dropbox.com/s/8uthfqkjarsbbre/mixed\\_trust\\_scheduling-tr.pdf?dl=0](https://www.dropbox.com/s/8uthfqkjarsbbre/mixed_trust_scheduling-tr.pdf?dl=0), 2018.
- 762
- 763 2 RTCA Special Committee 205. Formal methods supplement to DO-178C and DO-278A, 2011.
- 764 3 B. Andersson, S. Chaki, and D. de Niz. Combining symbolic runtime enforcers for cyber-  
765 physical systems. In *RV*, 2017.
- 766 4 E. Armbrust, J. Song, G. Bloom, and G. Parmer. On spatial isolation for mixed criticality,  
767 embedded systems. In *WMC*, 2014.
- 768 5 S. K. Baruah. Dynamic- and static-priority scheduling of recurring real-time tasks. *Journal of*  
769 *Real-Time Systems*, 2003.
- 770 6 S. Chaki and D. de Niz. Formal verification of a timing enforcer implementation. *ACM TECS*,  
771 2017.
- 772 7 M. C. Consiglio, J. P. Chamberlain, C. A. Munoz, and K. D. Hoffer. Concepts of integration  
773 for UAS operations in the NAS. In *Congress of the International Council of the Aeronautical*  
774 *Sciences (ICAS)*, 2012.
- 775 8 M. Correia, P. Verissimo, and N.F. Neves. The design of a COTS real-time distributed security  
776 kernel. In *EDCC*, 2002.
- 777 9 A. Crespo, I. Ripoll, and M. Masmano. Partitioned embedded architecture based on hypervisor:  
778 The XtratuM approach. In *EDCC*, 2010.
- 779 10 R. Davis and A. Burns. Mixed-criticality systems—a review. In *Technical Report, University*  
780 *of York, Available at <https://www-users.cs.york.ac.uk/burns/review.pdf>*, 2018.



© Copyright 2019 Carnegie Mellon University, Hyoseung Kim and John P. Lehoczky. All Rights Reserved.;

licensed under Creative Commons License CC-BY



Leibniz International Proceedings in Informatics

LIPIC Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

- 781 **11** R.I. Davis, A. Burns, R.J. Bril, and J.J. Lukkien. Controller area network (CAN) schedulability  
782 analysis: Refuted, revisited and revised. *Real-Time Systems*, 2007.
- 783 **12** D. de Niz, B. Andersson, and G. Moreno. Safety enforcement for the verification of autonomous  
784 systems. In *Proceedings of SPIE*, 2018.
- 785 **13** D. de Niz, K. Lakshmanan, and R. Rajkumar. On the scheduling of mixed-criticality real-time  
786 task sets. In *RTSS*, 2009.
- 787 **14** H. Ding, L. Arber, L. Sha, and M. Caccamo. The dependency management framework: a case  
788 study of the ION cubesat. In *ECRTS*, 2006.
- 789 **15** A. Easwaran, I. Lee, I. Shin, and O. Sokolsky. Compositional schedulability analysis of  
790 hierarchical real-time systems. In *ISORC*, 2007.
- 791 **16** R. Gu, Z. Shao, H. Chen, X. Wu, J. Kim, V. Sjöberg, and D. Costanzo. CertiKOS: An  
792 extensible architecture for building certified concurrent OS kernels. In *OSDI*, 2016.
- 793 **17** Z. Jiang, N.C. Audsley, and P. Dong. Bluevisor: A scalable real-time hardware hypervisor for  
794 many-core embedded systems. In *RTAS*, 2018.
- 795 **18** R. Kaiser and S. Wagner. Evolution of the PikeOS microkernel. In *First International  
796 Workshop on Microkernels for Embedded Systems*, 2007.
- 797 **19** G. Klein, K. Elphinstone, G. Heiser, J. Andronick, D. Cock, P. Derrin, D. Elkaduwe, K. Engel-  
798 hardt, R. Kolanski, M. Norrish, T. Sewell, H. Tuch, and S. Winwood. seL4: Formal verification  
799 of an OS kernel. In *SOSP*, 2009.
- 800 **20** J.P. Lehoczky. Fixed priority scheduling of periodic task sets with arbitrary deadlines. In  
801 *RTSS*, 1990.
- 802 **21** Y. Li, R. West, Z. Cheng, and E. Missimer. Predictable communication and migration in  
803 the Quest-V separation kernel. In *RTSS*, 2014.
- 804 **22** S. Mohan, M.-K. Yoon, R. Pellizzoni, and R. Bobba. Real-time systems security through  
805 scheduler constraints. In *ECRTS*, 2014.
- 806 **23** Special C. of RTCA. DO-178C, software considerations in airborne systems and equipment  
807 certification, 2011.
- 808 **24** D. Seto, J. P. Lehoczky, L. Sha, and K. G. Shin. On task schedulability in real-time  
809 control systems. In *17th IEEE Real-Time Systems Symposium*, pages 13–21, Dec 1996.  
810 doi:10.1109/REAL.1996.563693.
- 811 **25** L. Sha. Using simplicity to control complexity. *IEEE Software*, 2001.
- 812 **26** L. Singaravelu, C. Pu, H. Härtig, and C. Helmuth. Reducing TCB complexity for security-  
813 sensitive applications: Three case studies. In *Eurosys*, 2006.
- 814 **27** A. Vasudevan and S. Chaki. Have your PI and eat it too: Practical security on a low-cost  
815 ubiquitous computing platform. In *IEEE Euro Symposium on Security and Privacy*, 2018.
- 816 **28** A. Vasudevan, S. Chaki, L. Jia, J. M. McCune, J. Newsome, and A. Datta. Design, imple-  
817 mentation and verification of an eXtensible and Modular Hypervisor Framework. In *2013  
818 IEEE Symposium on Security and Privacy, SP*, 2013.
- 819 **29** A. Vasudevan, S. Chaki, P. Maniatis, L. Jia, and A. Datta. überspark: Enforcing verifiable  
820 object abstractions for automated compositional security analysis of a hypervisor. In *25th  
821 USENIX Security Symposium (USENIX Security 16)*, 2016.
- 822 **30** S. Xia, J. Wilson, C. Lu, and C.D. Gill. RT-Xen: Towards real-time hypervisor scheduling in  
823 Xen. In *EMSOFT*, 2011.

