

Exact Analysis of ZSRM Mixed-Criticality Scheduling of Sporadic Tasks

Björn Andersson, Dionisio de Niz, Hyoseung Kim, Mark Klein, and Rangunathan (Raj) Rajkumar
Carnegie Mellon University

Abstract—Zero-Slack Rate-Monotonic (ZSRM) is a family of mixed-criticality schedulers which are based on fixed-priority preemptive scheduling. One scheduler (which we call ZSRM-S) [5] works as follows: a job J is suspended at time t if at time t there is a higher-criticality job J' that has not finished and t minus the arrival time of J' exceeds a per-task configurable parameter (which we call zero-slack offset). ZSRM-S has two advantages compared to other mixed-criticality schedulers: (i) adaptation is local; i.e., there is no system-wide mode change needed and (ii) resumption is simple and natural. ZSRM-S has one drawback [7]: a high-criticality job J' can suffer from interference from a low-criticality job J that resumed after being suspended by another high-criticality job J'' (carry-in). Therefore, a variant of ZSRM (which we call ZSRM-SE) has been proposed [6]; it uses an enforcement mechanism to avoid carry-in. With ZSRM-SE, if a high-criticality job causes a low-criticality job to suspend and the high-criticality job has performed more execution than a certain bound then the low-criticality job shall not resume. We consider constrained-deadline sporadic tasks scheduled by ZSRM-S and present an exact schedulability test which solves a Mixed-Integer Linear Program (MILP). We also present that result for ZSRM-SE.

I. INTRODUCTION

The problem of scheduling real-time tasks with different criticalities is not new [9, 11] but the trend towards the increasing use of embedded computers and consolidating multiple functionalities onto a single computer platform has increased the importance of this problem. For this reason, researchers have, during recent years, developed more advanced schedulers and analysis methods for systems with tasks of different criticalities. The literature is extensive — see [3] for an excellent survey. Today, most schedulers for mixed criticality systems (MCS) rely on three ideas:

- I1.** A task is assigned a criticality level;
- I2.** If it is impossible to meet all deadlines and the scheduler has to let one task miss a deadline then the scheduler should let a lower criticality task miss a deadline;
- I3.** The execution time of a task is characterized by multiple numbers; each number is believed to be an upper bound on the execution time of the task but the confidence one has in this belief is different for different numbers.

The research community has used these ideas in different ways. One way is to extend schedulability analysis for classic fixed-priority or Earliest-Deadline-First so that when performing schedulability analysis to determine if task τ_i meets its deadlines then execution times of other tasks must be selected to be on the same confidence level as the criticality of task τ_i . It was found that many of the optimality results in non-MCS scheduling do not apply to MCS scheduling [12, 2]. Other

works use run-time monitoring and adaptation; check if a low criticality task has executed for more than it should and if so, the system switches to a high-critical mode where only high-critical tasks are allowed to execute [1]. Such an approach has two drawbacks: (i) it uses a system-wide mode and hence a system-wide mode-change is needed and (ii) it specifies how to switch from normal mode to an overload mode but typically does not specify how to switch back. We believe an alternative should be sought and hence, we consider the following idea:

I4. Before run-time, for each task τ_i , compute a parameter Z_i and at run-time, if a job of task τ_i has not finished at time Z_i after its arrival then take action to adapt.

This idea has been used for non-MCS and for this context, the action taken at Z_i is to change priorities; such use is called dual-priority scheduling [4]. This idea has been used for MCS and for this context, the action taken is to suspend jobs; such a scheduler was called ZSRM [5]. Later papers have discussed different semantics for it [7]; therefore, we let ZSRM denote a family of schedulers rather than a specific scheduler. ZSRM has been useful as witnessed by the following facts: previous work has made available implementations of a ZSRM scheduler in the Linux kernel and in VxWorks, as well as a sufficient schedulability test for it and this schedulability test is available to software practitioners in the OSATE AADL workbench. And a modification of it was used in a UAV system to ensure that an overload in vision processing does not jeopardize deadline guarantees of flight-control software [6]. Hence, we believe ZSRM is one of the most practical ideas in mixed-criticality scheduling. One scheduler (which we call ZSRM-S) [5] works as follows: a job J is suspended at time t if at time t there is a higher-criticality job J' that has not finished and t minus the arrival time of J' exceeds Z_i . Hence, the S in the name ZSRM-S means suspend. ZSRM-S has one drawback [7]: a high-criticality job J' can suffer from interference from a low-criticality job J that resumed after being suspended by another high-criticality job J'' (carry-in). Therefore, a variant of ZSRM (which we call ZSRM-SE) has been proposed [6]; it uses an enforcement mechanism to avoid carry-in. With ZSRM-SE, if a high-criticality job causes a low-criticality job to suspend and the high-criticality job has performed more execution than a certain bound then the low-criticality job shall not resume. Hence, the E in the name ZSRM-SE means execution-time monitoring. Unfortunately, no exact analysis was known for these schedulers.


Therefore, in this paper, we consider constrained-deadline sporadic tasks scheduled by ZSRM-S and present an exact

$$\begin{aligned}
|\tau| &= 2 \\
T_1 &= 4 \quad D_1 = 4 \quad C_1 = 2 \quad C_1^o = 2 \quad \zeta_1 = 1 \quad \text{prio}_1 = 2 \quad Z_1 = 2 \\
T_2 &= 10 \quad D_2 = 8 \quad C_2 = 2.5 \quad C_2^o = 5 \quad \zeta_2 = 2 \quad \text{prio}_2 = 1 \quad Z_2 = 5
\end{aligned}$$

Fig. 1: An example of a taskset in our model.

$$\begin{aligned}
\text{nj}_1(R) &= 2 \\
A_{1,1}(R) &= 0 \quad c_{1,1}(R) = 2 \quad A_{1,2}(R) = 4 \quad c_{1,2}(R) = 2 \\
\text{nj}_2(R) &= 1 \\
A_{2,1}(R) &= 0 \quad c_{2,1}(R) = 2.5
\end{aligned}$$

(a) An assignment R for the taskset in Fig. 1.



(b) ZSRM-S schedule of the assignment R for the taskset in Fig. 1.

Fig. 2: An assignment for the taskset in Fig. 1 and its ZSRM-S schedule. At time 5, the zero-slack instant of $\tau_{2,1}$ occurs so at this time, jobs from tasks with lower criticality get suspended; specifically, $\tau_{1,2}$ is suspended at this time. In (a), $c_{2,1}(R)$ is small so that when $\tau_{2,1}$ finishes, there is still time for $\tau_{1,2}$ to finish execution by its deadline (see (b)). For another assignment where $c_{2,1}$ is five, $\tau_{1,2}$ would miss its deadline.

schedulability test which solves a Mixed-Integer Linear Program (MILP). We also present that result for ZSRM-SE.

The rest of the paper is organized as follows. Section II presents the system model. Section III presents the new schedulability test for ZSRM-S. Section IV presents the new schedulability test for ZSRM-SE. Section V presents tools that perform the calculations of these schedulability tests. Section VI concludes.

II. SYSTEM MODEL

Throughout this paper, we let *s.t.* mean “such that” and we let $:$ mean “it holds that” and we let $\{x|f(x)\}$ denote a set of elements so that an element x is in the set if and only if $f(x)$ is true. We let $\langle a, b \rangle$ indicate a tuple with two elements a and b . We let $[a, b]$ indicate an interval of real numbers. We let $\{a..b\}$ indicate the set of integers that are $\geq a$ and $\leq b$.

Static parameters. We consider a system comprising a taskset τ and a computer platform comprising a single processor. A task τ_i in τ is characterized by T_i , D_i , C_i , C_i^o , ζ_i , prio_i , and Z_i with the interpretation that τ_i generates a sequence of jobs with two consecutive jobs of τ_i having arrival times separated by at least T_i and each job of τ_i must finish within D_i time units. ζ_i indicates the criticality of τ_i . (If ζ_i is high then the criticality of τ_i is high.) prio_i indicates the priority of τ_i . (If prio_i is high then the priority of τ_i is high.) We assume $\forall \tau_i \in \tau: C_i \leq C_i^o$ and $D_i \leq T_i$. The symbol Z_i means zero-slack offset of τ_i and it is used by the scheduler to determine the time instant when jobs of lower criticality should be adapted (e.g. suspended). The symbols C_i and C_i^o are upper bounds on the execution time of a job of τ_i ; the reason for having two upper bounds will be explained later in this section. For historical reasons, we refer to C_i as nominal execution time and C_i^o as overload execution time. Fig. 1 shows an example of a taskset in our model.

Run-time behavior of ZSRM-S. Let $\tau_{i,q}$ denote the q^{th} job of τ_i . Let R denote an assignment, for each task, the number of jobs it generates and for each of the jobs, an arrival time and execution time. Certain quantities that we define will be a function of the schedule and then we let *sc* be a schedule. Let $\text{nj}_i(R)$ denote the number of jobs that τ_i generates. Let $A_{i,q}(R)$ denote the arrival time of $\tau_{i,q}$. Let $c_{i,q}(R)$ denote the execution time of $\tau_{i,q}$. Let $f_{i,q}(\text{sc}, R)$ denote the finishing time of $\tau_{i,q}$. Let $\text{donex}_{i,q}(t, \text{sc}, R)$ denote the cumulative duration

of execution of $\tau_{i,q}$ before time t . Fig. 3 shows predicates that we use. $\text{elig}(i, q, t, \tau, R, \text{sc})$ is a predicate that is true if, at time t , the job $\tau_{i,q}$ has arrived but not finished and $\tau_{i,q}$ is not suspended at time t because of higher-criticality jobs. $\text{eligZSRMS}(i, q, t, \tau, R, \text{sc})$ indicates that $\tau_{i,q}$ is eligible for execution (i.e., it is in the ready queue or it is running) at time t . Clearly, because of priority-based scheduling, an eligible job will only execute if there is no other eligible job with higher priority. We use the predicate $\text{candZSRMS}(i, q, t, \tau, R, \text{sc})$ to indicate that $\tau_{i,q}$ is a candidate for execution; i.e., $\tau_{i,q}$ is eligible and there is no eligible job of higher priority. An instant is a ZSRMSschedinst if there is a job that arrives at this instant or there is a job that finishes at this instant or there is a job that has its zero-slack instant at this instant — see Fig. 3. At each instant t such that t is a ZSRMSschedinst, the scheduler does the following: if there is at least one job $\tau_{i,q}$ such that $\text{candZSRMS}(i, q, t, \tau, R, \text{sc})$, then arbitrarily choose a job $\tau_{i,q}$ such that $\text{candZSRMS}(i, q, t, \tau, R, \text{sc})$ and execute it on the processor at time t and let it continue to execute until the next ZSRMSschedinst; if there is no job $\tau_{i,q}$ such that $\text{candZSRMS}(i, q, t, \tau, R, \text{sc})$, then keep the processor idle at time t until the next schedinst. Fig. 2 shows a schedule that the taskset in Fig. 1 can generate.

Run-time behavior of ZSRM-SE. The run-time behavior of ZSRM-SE differs from ZSRM-S only in that with ZSRM-SE, a job J is terminated if there was a time now or in the past such that at that time, there was a higher-criticality job J' that has reached its zero-slack instant and not finished and executed for more than its nominal execution time. We specify this formally with predicates in Fig. 3. The predicate $\text{candZSRMSE}(i, q, t, \tau, R, \text{sc})$ indicates that $\tau_{i,q}$ is a candidate for execution. The predicate $\text{eligZSRMSE}(i, q, t, \tau, R, \text{sc})$ indicates that $\tau_{i,q}$ is eligible for execution; if $\text{terminatedZSRMSE}(i, q, t, \tau, R, \text{sc})$ is true then $\tau_{i,q}$ is not eligible for execution. The predicate $\text{terminatedZSRMSE}(i, q, t, \tau, R, \text{sc})$ is true if there is a time t' such that $t' \leq t$ and $\text{terminatednowZSRMSE}(i, q, t', \tau, R, \text{sc})$. The predicate $\text{terminatednowZSRMSE}(i, q, t', \tau, R, \text{sc})$ is true if there is job $\tau_{i',q'}$ such that $\zeta_{i'} > \zeta_i$ and $\tau_{i',q'}$ has arrived but not finished and $\tau_{i',q'}$ has executed for more than $C_{i'}$ time units.

Schedulability and schedulability test of ZSRM-S. We say that a job $\tau_{i,q}$ is success if its finishing time is at most

$$\begin{aligned}
\text{ZSRMSSchedinst}(t, \tau, R, \text{sc}) &= (\exists \langle i, q \rangle \text{ s.t. } (\tau_i \in \tau) \wedge (q \in 1..nj_i(R)) \wedge ((A_{i,q}(R) = t) \vee (f_{i,q}(R, \text{sc}) = t) \vee (A_{i,q}(R) + Z_i = t))) \\
\text{ZSRMSEschedinst}(t, \tau, R, \text{sc}) &= (\exists \langle i, q \rangle \text{ s.t. } (\tau_i \in \tau) \wedge (q \in 1..nj_i(R)) \wedge ((A_{i,q}(R) = t) \vee (f_{i,q}(R, \text{sc}) = t) \vee (A_{i,q}(R) + Z_i = t) \vee \\
&\quad (\text{donex}_{i',q'}(t, \text{sc}, R) = C_{i'}))) \\
\text{arrived}(i, q, t, \tau, R, \text{sc}) &= (A_{i,q}(R) \leq t) \\
\text{Zd}(i, q, t, \tau, R, \text{sc}) &= (A_{i,q}(R) + Z_i \leq t) \\
\text{finZSRMS}(i, q, t, \tau, R, \text{sc}) &= (f_{i,q}(\text{sc}, R) \leq t) \\
\text{finZSRMSE}(i, q, t, \tau, R, \text{sc}) &= ((f_{i,q}(\text{sc}, R) \leq t) \vee (\text{terminated}(i, q, t, \tau, R, \text{sc}))) \\
\text{arrivednotfinZSRMS}(i, q, t, \tau, R, \text{sc}) &= ((\text{arrived}(i, q, t, \tau, R, \text{sc})) \wedge (\neg \text{finZSRMS}(i, q, t, \tau, R, \text{sc}))) \\
\text{arrivednotfinZSRMSE}(i, q, t, \tau, R, \text{sc}) &= ((\text{arrived}(i, q, t, \tau, R, \text{sc})) \wedge (\neg \text{finZSRMSE}(i, q, t, \tau, R, \text{sc}))) \\
\text{ZdnotfinZSRMS}(i, q, t, \tau, R, \text{sc}) &= ((\text{Zd}(i, q, t, \tau, R, \text{sc})) \wedge (\neg \text{finZSRMS}(i, q, t, \tau, R, \text{sc}))) \\
\text{ZdnotfinZSRMSE}(i, q, t, \tau, R, \text{sc}) &= ((\text{Zd}(i, q, t, \tau, R, \text{sc})) \wedge (\neg \text{finZSRMSE}(i, q, t, \tau, R, \text{sc}))) \\
\text{lowexZSRMSE}(i, q, t, \tau, R, \text{sc}) &= (\text{donex}_{i,q}(t, \text{sc}, R) \leq C_i) \\
\text{ZdnotfinlowexZSRMSE}(i', q', t, \tau, R, \text{sc}) &= ((\text{ZdnotfinZSRMSE}(i, q, t, \tau, R, \text{sc})) \wedge (\text{lowexZSRMSE}(i, q, t, \tau, R, \text{sc}))) \\
\text{ZdnotfinhiexZSRMSE}(i', q', t, \tau, R, \text{sc}) &= ((\text{ZdnotfinZSRMSE}(i, q, t, \tau, R, \text{sc})) \wedge (\neg \text{lowexZSRMSE}(i, q, t, \tau, R, \text{sc}))) \\
\text{suspendednowZSRMS}(i, q, t, \tau, R, \text{sc}) &= (\exists \tau_{i',q'} \text{ s.t. } (\zeta_{i'} > \zeta_i) \wedge (\text{ZdnotfinZSRMS}(i', q', t, \tau, R, \text{sc}))) \\
\text{suspendednowZSRMSE}(i, q, t, \tau, R, \text{sc}) &= ((\exists \tau_{i',q'} \text{ s.t. } (\zeta_{i'} > \zeta_i) \wedge (\text{ZdnotfinZSRMSE}(i', q', t, \tau, R, \text{sc}))) \wedge \\
&\quad (\forall \tau_{i',q'} \text{ s.t. } (\zeta_{i'} > \zeta_i) : (\text{ZdnotfinZSRMSE}(i', q', t, \tau, R, \text{sc}) \Rightarrow \\
&\quad (\text{ZdnotfinlowexZSRMSE}(i', q', t, \tau, R, \text{sc})))))) \\
\text{terminatednowZSRMSE}(i, q, t, \tau, R, \text{sc}) &= (\exists \tau_{i',q'} \text{ s.t. } (\zeta_{i'} > \zeta_i) \wedge (\text{ZdnotfinhiexZSRMSE}(i', q', t, \tau, R, \text{sc}))) \\
\text{terminatedZSRMSE}(i, q, t, \tau, R, \text{sc}) &= (\exists t' \text{ s.t. } (t' \leq t) \wedge (\text{terminatednowZSRMSE}(i, q, t', \tau, R, \text{sc}))) \\
\text{eligZSRMS}(i, q, t, \tau, R, \text{sc}) &= ((\text{arrivednotfinZSRMS}(i, q, t, \tau, R, \text{sc})) \wedge (\neg \text{suspendednowZSRMS}(i, q, t, \tau, R, \text{sc}))) \\
\text{eligZSRMSE}(i, q, t, \tau, R, \text{sc}) &= ((\text{arrivednotfinZSRMSE}(i, q, t, \tau, R, \text{sc})) \wedge (\neg \text{suspendednowZSRMSE}(i, q, t, \tau, R, \text{sc}))) \wedge \\
&\quad (\neg \text{terminated}(i, q, t, \tau, R, \text{sc}))) \\
\text{candZSRMS}(i, q, t, \tau, R, \text{sc}) &= ((\text{eligZSRMS}(i, q, t, \tau, R, \text{sc})) \wedge (\forall \tau_{i',q'} \text{ s.t. } \text{prio}_{i'} > \text{prio}_i : \neg \text{eligZSRMS}(i', q', t, \tau, R, \text{sc}))) \\
\text{candZSRMSE}(i, q, t, \tau, R, \text{sc}) &= ((\text{eligZSRMSE}(i, q, t, \tau, R, \text{sc})) \wedge (\forall \tau_{i',q'} \text{ s.t. } \text{prio}_{i'} > \text{prio}_i : \neg \text{eligZSRMSE}(i', q', t, \tau, R, \text{sc}))) \\
\text{legMCS}(R, \text{sc}, \tau, \text{iD}, \text{qD}) &= ((\forall \langle i, q \rangle \text{ s.t. } (\tau_i \in \tau) \wedge (q \in 2..nj_i(R)) : A_{i,q}(R) - A_{i,q-1}(R) \geq T_i) \wedge \\
&\quad (\forall \langle i, q \rangle \text{ s.t. } (\tau_i \in \tau) \wedge (q \in 1..nj_i(R)) \wedge (\zeta_i > \zeta_{\text{iD}}) : c_{i,q}(R) \in [0, C_i]) \wedge \\
&\quad (\forall \langle i, q \rangle \text{ s.t. } (\tau_i \in \tau) \wedge (q \in 1..nj_i(R)) \wedge (\zeta_i \leq \zeta_{\text{iD}}) : c_{i,q}(R) \in [0, C_i^c])) \\
\text{successZSRMS}(i, q, \tau, R, \text{sc}) &= (f_{i,q}(\text{sc}, R) \leq A_{i,q}(R) + D_i) \\
\text{successZSRMSE}(i, q, \tau, R, \text{sc}) &= ((\neg \text{terminated}(i, q, A_{i,q}(R) + D_i, \tau, R, \text{sc})) \wedge (f_{i,q}(\text{sc}, R) \leq A_{i,q}(R) + D_i)) \\
\text{ZSRMSSch}(\tau) &= (\forall \langle \text{iD}, \text{qD}, R, \text{sc} \rangle \text{ s.t. } (\tau_{\text{iD}} \in \tau) \wedge (\text{qD} \in \{1..nj_{\text{iD}}(R)\}) \wedge \\
&\quad (\text{legMCS}(R, \text{sc}, \tau, \text{iD}, \text{qD})) \wedge (\text{legZSRMSSch}(\text{sc}, R, \tau)) : \text{successZSRMS}(\text{iD}, \text{qD}, \tau, R, \text{sc})) \\
\text{ZSRMSEsch}(\tau) &= (\forall \langle \text{iD}, \text{qD}, R, \text{sc} \rangle \text{ s.t. } (\tau_{\text{iD}} \in \tau) \wedge (\text{qD} \in \{1..nj_{\text{iD}}(R)\}) \wedge \\
&\quad (\text{legMCS}(R, \text{sc}, \tau, \text{iD}, \text{qD})) \wedge (\text{legZSRMSEsch}(\text{sc}, R, \tau)) : \text{successZSRMSE}(\text{iD}, \text{qD}, \tau, R, \text{sc}))
\end{aligned}$$

Fig. 3: Predicates that we will use.

its deadline. The predicate $\text{successZSRMS}(i, q, \tau, R, \text{sc})$ indicates that. The predicate $\text{legMCS}(R, \text{sc}, \tau, \text{iD}, \text{qD})$ is true if R satisfies certain constraints (expressing arrival times and execution times of jobs) — see Fig. 3. Note that compared to the definition of a legal assignment used in classic fixed-priority scheduling without mixed criticalities, our definition of legal assignment differs in two ways (i) our definition takes a schedule as input whereas the classic definition does not and (ii) our definition takes a task and job index as input whereas the classic does not. The predicate $\text{legZSRMSSch}(\text{sc}, R, \tau)$ indicates that schedule sc can be generated by ZSRM-S for assignment R for taskset τ . The predicate $\text{ZSRMSSch}(\tau)$ indicates that for each $\langle R, \text{iD}, \text{qD}, \text{sc} \rangle$ such that $\text{legMCS}(R, \text{sc}, \tau, \text{iD}, \text{qD})$ and $\text{legZSRMSSch}(\text{sc}, R, \tau)$, it holds that $\text{successZSRMS}(\text{iD}, \text{qD}, \tau, R, \text{sc})$. Intuitively, the meaning of $\text{ZSRMSSch}(\tau)$ is that $\text{ZSRMSSch}(\tau)$ is true if

each job J meets its deadline for the case that jobs of higher criticality than J have execution times that are bounded by the nominal execution times (not overload execution times). If $\text{ZSRMSSch}(\tau)$ is true then we say that the taskset is schedulable. Conversely, if $\text{ZSRMSSch}(\tau)$ is false then we say that the taskset is unschedulable. A schedulability test for ZSRM-S is a function that takes τ as input and outputs a boolean. For schedulability test ST associated with ZSRM-S, we say that ST is an *exact* schedulability test if $ST(\tau) \Leftrightarrow \text{ZSRMSSch}(\tau)$.

Schedulability and schedulability test of ZSRM-SE. The concepts for ZSRM-SE are analogous.

III. NEW SCHEDULABILITY TEST FOR ZSRM-S

Our goal in this section is to present an exact schedulability test for ZSRM-S. Traditional analysis of fixed-priority preemptive scheduling on a single processor relies on a con-

Sets :

$$\begin{aligned} \text{TS} &= \{i' | (\tau_{i'} \in \tau) \wedge ((\text{prio}_{i'} \geq \text{prio}_{\text{iD}}) \vee (\zeta_{i'} \geq \zeta_{\text{iD}}))\}, \text{TSHC}(i) = \{i' | (\tau_{i'} \in \tau) \wedge (\zeta_{i'} > \zeta_i)\}, \\ \text{TSHP}(i) &= \{i' | (\tau_{i'} \in \tau) \wedge (\text{prio}_{i'} > \text{prio}_i)\}, \text{QS}(i') = \{1.. \lceil \frac{t}{T_{i'}} \rceil\}, \text{PS} = \{1..3 \times \sum_{i' \in \text{TS}} \lceil \frac{t}{T_{i'}} \rceil\} \end{aligned}$$

Constraints :

$$\begin{aligned} t^1 &= 0 & \forall p \in (\text{PS} \setminus \{\text{PS}\}) : t^p &\leq t^{p+1} \\ \forall \langle i, q \rangle \text{ s.t. } (i \in \text{TS}) \wedge (q \in \text{QS}(i)) : & & \forall \langle i, q \rangle \text{ s.t. } (i \in \text{TS}) \wedge (q \in (\text{QS}(i) \setminus \{1\})) : A_{i,q} - A_{i,q-1} &\geq t_i \\ \forall \langle i, q, p \rangle \text{ s.t. } (i \in \text{TS}) \wedge (q \in \text{QS}(i)) \wedge (p \in \text{PS}) : & & \sum_{p' \in \text{PS}} \text{arrives}_{i,q}^{p'} = 1 & \sum_{p' \in \text{PS}} \text{finishes}_{i,q}^{p'} = 1 & \sum_{p' \in \text{PS}} \text{ZS}_{i,q}^{p'} = 1 \\ (\text{arrives}_{i,q}^p = 1) \Rightarrow (\text{donex}_{i,q}^p = 0) & (\text{arrives}_{i,q}^p = 1) \Rightarrow (A_{i,q} = t^p) & \text{arrived}_{i,q}^p = \sum_{p' \in \{1..p\}} \text{arrives}_{i,q}^{p'} \\ (\text{finishes}_{i,q}^p = 1) \Rightarrow (\text{donex}_{i,q}^p = c_{i,q}) & (\text{finishes}_{i,q}^p = 1) \Rightarrow (f_{i,q} = t^p) & \text{finZSRMS}_{i,q}^p = \sum_{p' \in \{1..p\}} \text{finishes}_{i,q}^{p'} \\ (\text{ZS}_{i,q}^p = 1) \Rightarrow (A_{i,q} + Z_i = t^p) & \text{Zd}_{i,q}^p = \sum_{p' \in \{1..p\}} \text{ZS}_{i,q}^{p'} \\ (\text{arrivednotfinZSRMS}_{i,q}^p = 1) \Leftrightarrow ((\text{arrived}_{i,q}^p = 1) \wedge (\text{finZSRMS}_{i,q}^p = 0)) \\ (\text{ZdnotfinZSRMS}_{i,q}^p = 1) \Leftrightarrow ((\text{Zd}_{i,q}^p = 1) \wedge (\text{finZSRMS}_{i,q}^p = 0)) \\ (\text{suspendednowZSRMS}_{i,q}^p = 1) \Rightarrow (\sum_{i' \in \text{TSHC}(i)} \sum_{q' \in \text{QS}(i')} \text{ZdnotfinZSRMS}_{i',q'}^p \geq 1) \\ (\text{suspendednowZSRMS}_{i,q}^p = 0) \Rightarrow (\sum_{i' \in \text{TSHC}(i)} \sum_{q' \in \text{QS}(i')} \text{ZdnotfinZSRMS}_{i',q'}^p \leq 0) \\ (\text{eligZSRMS}_{i,q}^p = 1) \Leftrightarrow ((\text{arrivednotfinZSRMS}_{i,q}^p = 1) \wedge (\text{suspendednowZSRMS}_{i,q}^p = 0)) \\ (\text{candZSRMS}_{i,q}^p = 1) \Leftrightarrow ((\text{eligZSRMS}_{i,q}^p = 1) \wedge (\wedge_{i' \in \text{TSHP}(i)} \wedge_{q' \in \text{QS}(i')} (\text{eligZSRMS}_{i',q'}^p = 0))) \\ \forall p \in (\text{PS} \setminus \{\text{PS}\}) : & & (\text{busyZSRMS}^p = 1) \Rightarrow (\sum_{\tau_{i'} \in \tau} \sum_{q' \in \text{QS}(i')} \text{candZSRMS}_{i',q'}^p \geq 1) & (\text{busyZSRMS}^p = 0) \Rightarrow (\sum_{\tau_{i'} \in \tau} \sum_{q' \in \text{QS}(i')} \text{candZSRMS}_{i',q'}^p \leq 0) \\ \forall \langle i, q, p \rangle \text{ s.t. } (i \in \text{TS}) \wedge (q \in \text{QS}(i)) \wedge (\text{PS} \setminus \{\text{PS}\}) : & & (\text{xZSRMS}_{i,q}^p = 1) \Rightarrow (\text{donex}_{i,q}^{p+1} = \text{donex}_{i,q}^p + t^{p+1} - t^p) & \text{x}_{i,q}^p \leq \text{candZSRMS}_{i,q}^p \\ \forall p \in (\text{PS} \setminus \{\text{PS}\}) : & & \sum_{\tau_{i'} \in \text{TS}} \sum_{q' \in \text{QS}(i')} \text{xZSRMS}_{i',q'}^p = \text{busyZSRMS}^p \\ \forall p \in \text{PS} : & & (\text{finishes}_{\text{iD},\text{qD}}^p = 1) \Rightarrow (\sum_{p' \in \{1..p-1\}} \text{busyZSRMS}^{p'} \geq p - 1) \\ \forall \langle i, q \rangle \text{ s.t. } (i \in \text{TS}) \wedge (q \in \text{QS}(i)) \wedge (\zeta_i > \zeta_{\text{iD}}) : & & c_{i,q} \leq C_i \\ \forall \langle i, q \rangle \text{ s.t. } (i \in \text{TS}) \wedge (q \in \text{QS}(i)) \wedge (\zeta_i \leq \zeta_{\text{iD}}) : & & c_{i,q} \leq C_i^o \end{aligned}$$

Domains of variables : $t^p \in \mathbb{R}_{\geq 0}$, $A_{i,q} \in \mathbb{R}_{\geq 0}$, $c_{i,q} \in \mathbb{R}_{\geq 0}$, $f_{i,q} \in \mathbb{R}_{\geq 0}$, $\text{donex}_{i,q}^p \in \mathbb{R}_{\geq 0}$, $\text{arrives}_{i,q}^p \in \{0, 1\}$, $\text{arrived}_{i,q}^p \in \{0, 1\}$, $\text{finishes}_{i,q}^p \in \{0, 1\}$, $\text{finZSRMS}_{i,q}^p \in \{0, 1\}$, $\text{ZS}_{i,q}^p \in \{0, 1\}$, $\text{Zd}_{i,q}^p \in \{0, 1\}$, $\text{arrivednotfinZSRMS}_{i,q}^p \in \{0, 1\}$, $\text{ZdnotfinZSRMS}_{i,q}^p \in \{0, 1\}$, $\text{suspendednowZSRMS}_{i,q}^p \in \{0, 1\}$, $\text{eligZSRMS}_{i,q}^p \in \{0, 1\}$, $\text{candZSRMS}_{i,q}^p \in \{0, 1\}$, $\text{busyZSRMS}_{i,q}^p \in \{0, 1\}$, $\text{xZSRMS}_{i,q}^p \in \{0, 1\}$

Fig. 4: Constraints we use for exact schedulability analysis of ZSRM-S.

dition for critical instant [10] or the concept busy period [8]. Unfortunately, these concepts cannot be used directly for exact schedulability analysis of ZSRM-S because in ZSRM-S, a job may be suspended. Thus, we will develop new ideas and put them together into an exact schedulability test for ZSRM-S.

For this discussion, let $\text{feas}(X)$ denote a predicate that is true if and only if X (a set of constraints) is feasible. Let $\max\{\text{myobj} | X\}$ denote the largest value of myobj subject to the constraints X . Also, let $\text{ct}(t, \text{iD}, \text{qD})$ denote the set of constraints in Fig. 4 where (i) t in Fig. 4 is a constant which is equal to the 1st parameter of ct , (ii) iD in Fig. 4 is a constant

which is equal to the 2nd parameter of ct , and (iii) qD in Fig. 4 is a constant which is equal to the 3rd parameter of ct .

Our first lemma states certain properties of a time interval for an unschedulable taskset.

Lemma 1.

$$\begin{aligned} (\neg \text{ZSRMSsch}(\tau)) \Rightarrow & \\ (\exists \text{iD}, \text{qD}, R, \text{sc}) \text{ s.t. } (\tau_{\text{iD}} \in \tau) \wedge (\text{qD} \in \{1..n_{\text{jID}}(R)\}) \wedge & \\ (\text{legMCS}(R, \text{sc}, \tau, \text{iD}, \text{qD})) \wedge (\text{legZSRMSsch}(\text{sc}, R, \tau)) \wedge & \\ (f_{\text{iD},\text{qD}}(\text{sc}, R) - A_{\text{iD},\text{qD}} > D_{\text{iD}}) \wedge & \end{aligned}$$

(for schedule sc , it holds that at all times
before time 0, the processor is idle) \wedge
(for schedule sc , it holds that at all times
in $[0, f_{iD,qD}(sc, R)]$, the processor executes
a job with priority $\geq \text{prio}_{iD}$ or criticality $\geq \zeta_{iD}$)

Proof: Assume that the left-hand side of the lemma is true. Then, from the definition of ZSRMSsch, it holds that:

$$\begin{aligned} & \exists \langle iD, qD, R, sc \rangle \text{ s.t. } (\tau_{iD} \in \tau) \wedge (qD \in \{1..n_{j_{iD}}(R)\}) \wedge \\ & (\text{legMCS}(R, sc, \tau, iD, qD)) \wedge (\text{legZSRMSsch}(sc, R, \tau)) \wedge \\ & (\neg \text{successZSRMS}(iD, qD, \tau, R, sc)) \end{aligned}$$

From the definition of successZSRMS we obtain an inequality, which applied on the above yields that:

$$\begin{aligned} & \exists \langle iD, qD, R, sc \rangle \text{ s.t. } (\tau_{iD} \in \tau) \wedge (qD \in \{1..n_{j_{iD}}(R)\}) \wedge \\ & (\text{legMCS}(R, sc, \tau, iD, qD)) \wedge (\text{legZSRMSsch}(sc, R, \tau)) \wedge \\ & (f_{iD,qD}(sc, R) - A_{iD,qD} > D_{iD}) \end{aligned} \quad (1)$$

In the schedule sc above, we can form a time interval that (i) ends at time $f_{iD,qD}(sc, R)$ and (ii) begins at the earliest time such that in this time interval, only jobs with priority $\geq \text{prio}_{iD}$ or criticality $\geq \zeta_{iD}$ execute. One can delete all jobs arriving before this time interval. We can also set the origin of the time axis to be such that time zero is the time when this time interval begins. Applying this reasoning on (1) yields:

$$\begin{aligned} & \exists \langle iD, qD, R, sc \rangle \text{ s.t. } (\tau_{iD} \in \tau) \wedge (qD \in \{1..n_{j_{iD}}(R)\}) \wedge \\ & (\text{legMCS}(R, sc, \tau, iD, qD)) \wedge (\text{legZSRMSsch}(sc, R, \tau)) \wedge \\ & (f_{iD,qD}(sc, R) - A_{iD,qD} > D_{iD}) \wedge \\ & (\text{for schedule } sc, \text{ it holds that at all times} \\ & \text{before time 0, the processor is idle}) \wedge \\ & (\text{for schedule } sc, \text{ it holds that at all times} \\ & \text{in } [0, f_{iD,qD}(sc, R)], \text{ the processor executes} \\ & \text{a job with priority } \geq \text{prio}_{iD} \text{ or criticality } \geq \zeta_{iD}) \end{aligned}$$

This is the right-hand side of the lemma. \blacksquare

Our second lemma states that if the taskset is unschedulable then there exists a tuple $\langle iD, qD, t \rangle$ such that a certain problem is infeasible and t is at most a certain bound. When expressing this bound, we need to compute

$$\max_{qD' \in \{1.. \lceil \frac{t}{T_{iD}} \rceil\} \wedge (\text{feas}(\text{ct}(t', iD, qD')))} \max\{f_{iD,qD'} | \text{ct}(t', iD, qD')\} \quad (2)$$

This expression means the following: (i) iterate over all qD' (which may be different from qD), (ii) check if the current qD' is such that $\text{ct}(t', iD, qD')$ is feasible, (iii) if the answer to the preceding question is yes, then evaluate $\max\{f_{iD,qD'} | \text{ct}(t', iD, qD')\}$, and (iv) take the maximum of the computed values. One can see that for $qD' = 1$, it holds that $\text{ct}(t', iD, qD')$ is feasible. Hence, the evaluation in step (ii) is true for at least one iteration. And hence, the expression in

(2) is well defined. With this, we can state our second lemma.

Lemma 2.

$$\begin{aligned} & (\neg \text{ZSRMSsch}(\tau)) \Rightarrow \\ & (\exists \langle iD, qD, t \rangle \text{ s.t. } (t > 0) \wedge (\tau_{iD} \in \tau) \wedge (qD \in \{1.. \lceil \frac{t}{T_{iD}} \rceil\}) \wedge \\ & (\text{feas}(\{f_{iD,qD} - A_{iD,qD} > D_{iD}\} \cup \text{ct}(t, iD, qD))) \wedge \\ & (t \leq \\ & \min\{t' | t' = \max_{qD' \in \{1.. \lceil \frac{t'}{T_{iD}} \rceil\} \wedge (\text{feas}(\text{ct}(t', iD, qD')))} \\ & \max\{f_{iD,qD'} | \text{ct}(t', iD, qD')\})\}) \end{aligned}$$

Proof: Assume that the left-hand side of the lemma is true. Then, using Lemma 1 yields that:

$$\begin{aligned} & \exists \langle iD, qD, R, sc \rangle \text{ s.t. } (\tau_{iD} \in \tau) \wedge (qD \in \{1..n_{j_{iD}}(R)\}) \wedge \\ & (\text{legMCS}(R, sc, \tau, iD, qD)) \wedge (\text{legZSRMSsch}(sc, R, \tau)) \wedge \\ & (f_{iD,qD}(sc, R) - A_{iD,qD} > D_{iD}) \wedge \\ & (\text{for schedule } sc, \text{ it holds that at all times} \\ & \text{before time 0, the processor is idle}) \wedge \\ & (\text{for schedule } sc, \text{ it holds that at all times} \\ & \text{in } [0, f_{iD,qD}(sc, R)], \text{ the processor executes} \\ & \text{a job with priority } \geq \text{prio}_{iD} \text{ or criticality } \geq \zeta_{iD}) \end{aligned} \quad (3)$$

Clearly, in the schedule sc above, the rules of dispatching (expressed in Fig. 3) applies and the assignment R is legal. Let us consider the part of schedule sc during $[0, f_{iD,qD}(sc, R)]$ and let us introduce t as $t = f_{iD,qD}(sc, R)$. We can encode this schedule with variables and constraints — indeed $\text{ct}(t, iD, qD)$, expressed in Fig. 4 does that. One can understand this encoding as follows: Clearly, for each task $\tau_{i'}$, there are at most $\lceil \frac{t}{T_{i'}} \rceil$ jobs of $\tau_{i'}$. Then we introduce variables that are direct analogs of the assignment R . The variable $A_{i',q'}$ in Fig. 4 is the arrival time of $\tau_{i',q'}$ and $c_{i',q'}$ in Fig. 4 is the execution time of $\tau_{i',q'}$. In Fig. 4, TS denotes the set of task that can generate jobs that can execute in the time interval $[0, t]$. In Fig. 4, QS(i') denotes the set of indices of jobs of task $\tau_{i'}$ that can execute in the time interval $[0, t]$. Recall that an instant is a schedinst if there is a job that arrives at this instant or there is a job that finishes at this instant or there is a job that has its zero-slack instant at this instant. Since we consider a time interval of duration t , and no jobs arrive before the time interval, it holds that there are at most $3 \times \sum_{i' \in \text{TS}} \lceil \frac{t}{T_{i'}} \rceil$ instants that are schedinst. We can divide time into sub-time-intervals that are non-intersecting and that these instants separate the sub-time-intervals. This gives us that (i) a sub-time-interval begins at an instant that is a schedinst and (ii) if a sub-time-interval is not the last one, then it ends at an instant that is a schedinst and (iii) within a sub-time-interval, there is no instant that is a schedinst. We call these sub-time-intervals *positions* and we let t^p denote the time when the p^{th} position starts. There are at most $3 \times \sum_{i' \in \text{TS}} \lceil \frac{t}{T_{i'}} \rceil - 1$ positions. We let PS denote the set $\text{PS} = \{1..3 \times \sum_{i' \in \text{TS}} \lceil \frac{t}{T_{i'}} \rceil\}$, i.e., if

$p' \leq |\text{PS}| - 1$ then $t^{p'}$ is the beginning of the p'^{th} position and if $p' = |\text{PS}|$ then $t^{p'}$ is the end of the $p' - 1^{\text{th}}$ position (the last position). Since we consider the time interval $[0, t]$, it holds that the first position starts at time 0, that is, $t^1 = 0$. For each position that is not the first position, it holds that its starting time is constrained to be at least as large as the starting time of its predecessor position; we express it as $\forall p \in (\text{PS} \setminus \{1\}) : t^{p-1} \leq t^p$.

We can then express whether an event occurs in the beginning of a position. $\text{arrives}_{i,q}^p$ is a variable in $\{0, 1\}$; if $\text{arrives}_{i,q}^p = 1$ then it means that $\tau_{i,q}$ arrives in the beginning of position p . $\text{arrived}_{i,q}^p$ is a variable in $\{0, 1\}$; if $\text{arrived}_{i,q}^p = 1$ then it means that $\tau_{i,q}$ arrives in the beginning of position p or in an earlier position. $\text{finishes}_{i,q}^p$ is a variable in $\{0, 1\}$; if $\text{finishes}_{i,q}^p = 1$ then it means that $\tau_{i,q}$ finishes in the beginning of position p . $\text{finZSRMS}_{i,q}^p$ is a variable in $\{0, 1\}$; if $\text{finZSRMS}_{i,q}^p = 1$ then it means that $\tau_{i,q}$ finishes in the beginning of position p or in an earlier position. $\text{ZSRMS}_{i,q}^p$ is a variable in $\{0, 1\}$; if $\text{ZSRMS}_{i,q}^p = 1$ then it means that $\tau_{i,q}$ arrives exactly Z_i before the beginning of position p . $\text{Zd}_{i,q}^p$ is a variable in $\{0, 1\}$; if $\text{Zd}_{i,q}^p = 1$ then it means that there is a position $p' \leq p$ such that $\text{ZSRMS}_{i,q}^{p'} = 1$. Fig. 3 shows predicates and these predicates describe dispatching. We can introduce variables that describe if a predicate is true for a job at a time which is the beginning of a position. For example, $\text{arrivednotfinZSRMS}_{i,q}^p$ is a variable in $\{0, 1\}$; if $\text{arrivednotfinZSRMS}_{i,q}^p = 1$ then it means that $\tau_{i,q}$ arrives in the beginning of position p or earlier and $\tau_{i,q}$ finishes in the beginning of a position later than p . In Fig. 4, we express this as $(\text{arrivednotfinZSRMS}_{i,q}^p = 1) \Leftrightarrow ((\text{arrived}_{i,q}^p = 1) \wedge (\text{finZSRMS}_{i,q}^p = 0))$. Other variables in Fig. 4 describe predicates in Fig. 3 analogously. In the end, we obtain a variable $\text{candZSRMS}_{i,q}^p$ which describes that $\tau_{i,q}$ is a candidate for execution in the beginning of position p . Recall that a job is a candidate for execution if it is eligible and there is no other eligible job with higher priority at this time. We then introduce $\text{xZSRMS}_{i,q}^p$ which is a variable in $\{0, 1\}$; if $\text{xZSRMS}_{i,q}^p = 1$ then it means that $\tau_{i,q}$ executes in position p . Clearly, a job can only execute if it is a candidate. In Fig. 4, we express this as $\text{xZSRMS}_{i,q}^p \leq \text{candZSRMS}_{i,q}^p$. The above reasoning yields:

$$\begin{aligned} \exists \langle iD, qD, t \rangle \text{ s.t. } (t > 0) \wedge (\tau_{iD} \in \tau) \wedge (qD \in \{1.. \lceil \frac{t}{T_{iD}} \rceil\}) \wedge \\ (\text{feas}(\{f_{iD,qD} - A_{iD,qD} > D_{iD}\} \cup \{f_{iD,qD} = t\} \cup \\ \text{ct}(t, iD, qD))) \end{aligned} \quad (4)$$

Let us now discuss the length of the busy period mentioned in (3). It can be seen that if $\tau_{iD,qD}$ misses its deadline in a time interval where only jobs with priority $\geq \text{prio}_{iD}$ or criticality $\geq \zeta_{iD}$, then

$$f_{iD,qD}(\text{sc}, R) \leq \min\{t' | t' = \max\{f_{iD,qD} | \text{ct}(t', iD, qD)\}\}$$

Clearly, since $t = f_{iD,qD}(\text{sc}, R)$, we obtain:

$$t \leq \min\{t' | t' = \max\{f_{iD,qD} | \text{ct}(t', iD, qD)\}\}$$

The right-hand side of this expression contains the symbol qD . We would like to find an upper bound that does not depend on qD . It can be seen that:

$$t \leq \min\{t' | t' = \max_{qD' \in \{1.. \lceil \frac{t'}{T_{iD}} \rceil\}} \max\{f_{iD,qD'} | \text{ct}(t', iD, qD')\}\}$$

Combining it with (4) yields:

$$\begin{aligned} \exists \langle iD, qD, t \rangle \text{ s.t. } (t > 0) \wedge (\tau_{iD} \in \tau) \wedge (qD \in \{1.. \lceil \frac{t}{T_{iD}} \rceil\}) \wedge \\ (\text{feas}(\{f_{iD,qD} - A_{iD,qD} > D_{iD}\} \cup \{f_{iD,qD} = t\} \cup \\ \text{ct}(t, iD, qD))) \wedge \\ (t \leq \\ \min\{t' | t' = \max_{qD' \in \{1.. \lceil \frac{t'}{T_{iD}} \rceil\} \wedge (\text{feas}(\text{ct}(t', iD, qD')))} \\ \max\{f_{iD,qD'} | \text{ct}(t', iD, qD')\}\}) \end{aligned}$$

Dropping one constraints cannot cause infeasibility. Hence, by dropping $\{f_{iD,qD} = t\}$ from the above, we have:

$$\begin{aligned} \exists \langle iD, qD, t \rangle \text{ s.t. } (t > 0) \wedge (\tau_{iD} \in \tau) \wedge (qD \in \{1.. \lceil \frac{t}{T_{iD}} \rceil\}) \wedge \\ (\text{feas}(\{f_{iD,qD} - A_{iD,qD} > D_{iD}\} \cup \text{ct}(t, iD, qD))) \wedge \\ (t \leq \\ \min\{t' | t' = \max_{qD' \in \{1.. \lceil \frac{t'}{T_{iD}} \rceil\} \wedge (\text{feas}(\text{ct}(t', iD, qD')))} \\ \max\{f_{iD,qD'} | \text{ct}(t', iD, qD')\}\}) \end{aligned}$$

This is the right-hand side of the lemma. ■

Our third lemma states how a change in t impacts certain inequalities.

Lemma 3. *If $t_a < t_b$ then it holds that:*

$$\begin{aligned} (\exists \langle iD, qD \rangle \text{ s.t. } (\tau_{iD} \in \tau) \wedge (qD \in \{1.. \lceil \frac{t_a}{T_{iD}} \rceil\}) \wedge \\ (\text{feas}(\{f_{iD,qD} - A_{iD,qD} > D_{iD}\} \cup \text{ct}(t_a, iD, qD)))) \\ \Rightarrow \\ (\exists \langle iD, qD \rangle \text{ s.t. } (\tau_{iD} \in \tau) \wedge (qD \in \{1.. \lceil \frac{t_b}{T_{iD}} \rceil\}) \wedge \\ (\text{feas}(\{f_{iD,qD} - A_{iD,qD} > D_{iD}\} \cup \text{ct}(t_b, iD, qD)))) \end{aligned}$$

Proof: Assume that the left-hand side is true. Since the left-hand side is true, we know that there is a solution to the constraints. We can copy that solution to use it to satisfy the constraints on the right-hand side and then for the new variables that only exists on the right-hand side but not on the left-hand side, we can set them to zero. This yields a solution to the constraints on the right-hand side. And hence the right-hand side is true. ■

Our fourth lemma states certain inequalities for an unschedulable taskset (it differs from the second lemma only in that it uses $=$ instead of \leq) on the right-hand side.

Lemma 4.

$$\begin{aligned}
& (\neg \text{ZSRMSsch}(\tau)) \Rightarrow \\
& (\exists \langle iD, qD, t \rangle \text{ s.t. } (t > 0) \wedge (\tau_{iD} \in \tau) \wedge (qD \in \{1.. \lceil \frac{t}{T_{iD}} \rceil\}) \wedge \\
& \quad (\text{feas}(\{f_{iD, qD} - A_{iD, qD} > D_{iD}\} \cup \text{ct}(t, iD, qD))) \wedge \\
& \quad (t = \\
& \quad \quad \min\{t' | t' = \max_{qD' \in \{1.. \lceil \frac{t'}{T_{iD}} \rceil\} \wedge (\text{feas}(\text{ct}(t', iD, qD')))} \\
& \quad \quad \quad \max\{f_{iD, qD'} | \text{ct}(t', iD, qD')\})\})
\end{aligned}$$

Proof: Follows from applying Lemma 3 on Lemma 2. ■

We will now discuss another direction of implication; we will discuss \Leftarrow instead of \Rightarrow . Our fifth lemma states that if certain inequalities are true then the taskset is unschedulable.

Lemma 5.

$$\begin{aligned}
& (\neg \text{ZSRMSsch}(\tau)) \Leftarrow \\
& (\exists \langle iD, qD, t \rangle \text{ s.t. } (t > 0) \wedge (\tau_{iD} \in \tau) \wedge (qD \in \{1.. \lceil \frac{t}{T_{iD}} \rceil\}) \wedge \\
& \quad (\text{feas}(\{f_{iD, qD} - A_{iD, qD} > D_{iD}\} \cup \text{ct}(t, iD, qD)))
\end{aligned}$$

Proof: Assume that the right-hand side of the lemma is true. Then there exists a $\langle iD, qD, t \rangle$ such that the right-hand side is true. Since the constraints on the right-hand side are feasible, we have an assignment of values to the variables in $\text{ct}(t, iD, qD)$ and with this assignment of values to variables we obtain an assignment R and can (based on the discussion in Lemma 2 and using R), construct a schedule during $[0, t]$ such that $\tau_{iD, qD}$ misses its deadline. Hence, it holds that

$$\begin{aligned}
& \exists \langle iD, qD, R, sc \rangle \text{ s.t. } (\tau_{iD} \in \tau) \wedge (qD \in \{1.. n_{j_{iD}}(R)\}) \wedge \\
& \quad (\text{legMCS}(R, sc, \tau, iD, qD)) \wedge (\text{legZSRMSsch}(sc, R, \tau)) \wedge \\
& \quad \quad (f_{iD, qD} - A_{iD, qD} > D_{iD})
\end{aligned}$$

This can be rewritten as :

$$(\neg \text{ZSRMSsch}(\tau))$$

This is the left-hand side of the lemma. ■

We can consider Lemma 5 but add additional constraints on the right-hand side.

Lemma 6.

$$\begin{aligned}
& (\neg \text{ZSRMSsch}(\tau)) \Leftarrow \\
& (\exists \langle iD, qD, t \rangle \text{ s.t. } (t > 0) \wedge (\tau_{iD} \in \tau) \wedge (qD \in \{1.. \lceil \frac{t}{T_{iD}} \rceil\}) \wedge \\
& \quad (\text{feas}(\{f_{iD, qD} - A_{iD, qD} > D_{iD}\} \cup \text{ct}(t, iD, qD))) \wedge \\
& \quad (t = \\
& \quad \quad \min\{t' | t' = \max_{qD' \in \{1.. \lceil \frac{t'}{T_{iD}} \rceil\} \wedge (\text{feas}(\text{ct}(t', iD, qD')))} \\
& \quad \quad \quad \max\{f_{iD, qD'} | \text{ct}(t', iD, qD')\})\})
\end{aligned}$$

Proof: Follows from Lemma 5. ■

We then present an exact condition for unschedulability.

1. allOK := true
2. **for each** $\tau_{iD} \in \tau$, as long as allOK **do**
3. $t := -1$; newt := C_{iD}^o
4. **while** ($t < \text{newt}$) **do**
5. $t := \text{newt}$
6. flag := false
7. **for each** $qD' \in \{1.. \lceil \frac{t}{T_{iD}} \rceil\}$ **do**
8. $\langle fe, va \rangle := \text{solve}(\max\{f_{iD, qD'} | \text{ct}(t, iD, qD')\})$
9. **if** fe **then**
10. **if** flag **then** newt := $\max(\text{newt}, va)$
11. **else** newt := va; flag := true
12. **end if**
13. **end if**
14. **end for**
15. **end while**
16. **for each** $qD \in \{1.. \lceil \frac{t}{T_{iD}} \rceil\}$, as long as allOK **do**
17. **if** $\text{feas}(\{f_{iD, qD} - A_{iD, qD} > D_{iD}\} \cup \text{ct}(t, iD, qD))$ **then**
18. allOK := false
19. **end if**
20. **end for**
21. **end for**
22. return allOK

Fig. 5: An algorithm for ZSRM-S schedulability testing.

Lemma 7.

$$\begin{aligned}
& (\neg \text{ZSRMSsch}(\tau)) \Leftrightarrow \\
& (\exists \langle iD, qD, t \rangle \text{ s.t. } (t > 0) \wedge (\tau_{iD} \in \tau) \wedge (qD \in \{1.. \lceil \frac{t}{T_{iD}} \rceil\}) \wedge \\
& \quad (\text{feas}(\{f_{iD, qD} - A_{iD, qD} > D_{iD}\} \cup \text{ct}(t, iD, qD))) \wedge \\
& \quad (t = \\
& \quad \quad \min\{t' | t' = \max_{qD' \in \{1.. \lceil \frac{t'}{T_{iD}} \rceil\} \wedge (\text{feas}(\text{ct}(t', iD, qD')))} \\
& \quad \quad \quad \max\{f_{iD, qD'} | \text{ct}(t', iD, qD')\})\})
\end{aligned}$$

Proof: Follows from Lemma 4 and Lemma 6 ■

We then present an exact conditions for schedulability.

Theorem 1.

$$\text{ZSRMSsch}(\tau) \Leftrightarrow$$

$$(\forall iD \text{ s.t. } (\tau_{iD} \in \tau) :$$

for $t =$

$$\min\{t' | t' = \max_{qD' \in \{1.. \lceil \frac{t'}{T_{iD}} \rceil\} \wedge (\text{feas}(\text{ct}(t', iD, qD')))} \max\{f_{iD, qD'} | \text{ct}(t', iD, qD')\}\} :$$

$$\forall qD \in \{1.. \lceil \frac{t}{T_{iD}} \rceil\} :$$

$$\neg \text{feas}(\{f_{iD, qD} - A_{iD, qD} > D_{iD}\} \cup \text{ct}(t, iD, qD))$$

Proof: Follows from rewriting Lemma 7. ■

Evaluating the right-hand side of Theorem 1 requires calculating $\min\{t' | t' = \max_{qD' \in \{1.. \lceil \frac{t'}{T_{iD}} \rceil\}} \max\{f_{iD, qD'} | \text{ct}(t', iD, qD')\}\}$ — let t_{\min} denote this. It can be seen that $\max_{qD' \in \{1.. \lceil \frac{t'}{T_{iD}} \rceil\}} \max\{f_{iD, qD'} | \text{ct}(t', iD, qD')\}$ is non-increasing with increasing t' (follows from reasoning

```

1. allOK := true
2. for each  $\tau_{iD} \in \tau$ , as long as allOK do
3.    $t := -1$ ; newt :=  $C_{iD}^o$ 
4.   while ( $t < \text{newt}$ ) and allOK do
5.      $t := \text{newt}$ 
6.     for each  $qD \in \{1.. \lceil \frac{t}{T_{iD}} \rceil\}$ , as long as allOK do
7.       if  $\text{feas}(\{f_{iD,qD} - A_{iD,qD} > D_{iD}\} \cup$ 
8.          $\text{ct}(t, iD, qD))$  then
9.         allOK := false
10.      end if
11.    end for
12.    if allOK then
13.      flag := false
14.      for each  $qD' \in \{1.. \lceil \frac{t}{T_{iD}} \rceil\}$  do
15.         $\langle fe, va \rangle := \text{solve}(\max\{f_{iD,qD'} | \text{ct}(t, iD, qD')\})$ 
16.        if fe then
17.          if flag then newt :=  $\max(\text{newt}, va)$ 
18.          else newt := va; flag := true
19.        end if
20.      end for
21.    end if
22.  end while
23. end for
24. return allOK

```

Fig. 6: An algorithm for ZSRM-S schedulability testing; it is optimized for detecting clearly unschedulable tasksets quickly.

analogous to the reasoning in the proof of Lemma 3). Therefore, we can evaluate t_{\min} with standard iterative procedure. Fig. 5 is an algorithm that uses such an iterative procedure to perform the schedulability test as expressed by Theorem 1. We use the notation $\langle fe, va \rangle := \text{solve}(\max\{\text{PROB}\})$ to state that the optimization problem PROB should be solved and fe is a boolean which is true if the problem is feasible and false otherwise; if the problem is feasible then va is the value of the objective function for an optimal solution.

For many tasksets that are unschedulable, it holds that there is a job that misses its deadline at an early time in a busy period. Unfortunately, Fig. 5 requires that we obtain $t \geq \text{newt}$ before we can even start checking the existence of deadline misses. We would like to get early termination for such tasksets. By using Lemma 5, we can check a given $\langle iD, qD, t \rangle$ to see if it satisfies certain conditions and if this is the case, we know that the taskset is unschedulable. We can apply this condition for the t after line 5 in Fig. 5. By adding such a check, we know that lines 16 to 20 in Fig. 5 are not needed. With these observations, we can rewrite the algorithm Fig. 5 into the algorithm in Fig. 6.

IV. NEW SCHEDULABILITY TEST FOR ZSRM-SE

In this section, we present an exact schedulability test for ZSRM-SE. Note that ZSRM-SE differs from ZSRM-S in only two ways. First, the definition of success of a job is different; a job can be not-success if the job misses its deadline (just like in ZSRM-S) but a job can also be not-success if it is terminated. Second, the schedules that can be generated by ZSRM-SE are different from the ones that can be generated

by ZSRM-S. Hence, use the constraints in Fig. 4 as a starting point and observe that ZSRM-SE is impacted by termination condition and hence, we add constraints for that and this results in the constraints in Fig. 7. Note that in Fig. 7, we have a variable $\text{terminatednow}_{i,q}^p$ with the interpretation that if $\text{terminatednow}_{i,q}^p = 1$ then $\tau_{i,q}$ is terminated at the beginning of position p . There is also a predicate $\text{terminated}_{i,q}^p$ with the interpretation that if $\text{terminated}_{i,q}^p = 1$ then $\tau_{i,q}$ is terminated at the beginning of position p or earlier. With these variables, we can define $\text{eligZSRMSE}_{i,q}^p$ that describe whether $\tau_{i,q}$ is eligible at the beginning of position p ; it is calculated based on $\text{terminated}_{i,q}^p$. Therefore, if $\tau_{iD,qD}$ is a not-success job then it holds that there is a t such that the following constraints are feasible: $\{((\text{terminated}_{iD,qD}^{\text{PS}} = 1) \vee (f_{iD,qD} > A_{iD,qD} + D_{iD}))\} \cup \text{ct2}(t, iD, qD)$, where ct2 is the set of constraints in Fig. 7 and PS is the last position. We also introduce $ft_{i,q}$ — meaning failure time — which is a variable that states the time that $\tau_{i,q}$ generated a failure. If $\tau_{i,q}$ is terminated then $ft_{i,q}$ is the time when it got terminated. If $\tau_{i,q}$ is not terminated then $ft_{i,q}$ is the time when it finished. Our formulation here also differ from the one in the previous section in that the number of scheduling instants is greater; here each job can generate four scheduling instants — the time when a job has executed exactly its nominal execution time can be a scheduling instant as well (because job termination can happen at such an instant).

Theorem 2.

$$\begin{aligned}
& \text{ZSRMSEsch}(\tau) \Leftrightarrow \\
& (\forall iD \text{ s.t. } (\tau_{iD} \in \tau) : \\
& \quad \text{for } t = \\
& \quad \quad \min\{t' | t' = \max_{qD' \in \{1.. \lceil \frac{t'}{T_{iD}} \rceil\}} \max\{ft_{iD,qD'} | \text{ct2}(t', iD, qD')\}\} : \\
& \quad \quad \quad \forall qD \in \{1.. \lceil \frac{t}{T_{iD}} \rceil\} : \\
& \quad \quad \quad \neg \text{feas}(\{(\text{terminated}_{iD,qD}^{\text{PS}} = 1) \vee (f_{iD,qD} - A_{iD,qD} > D_{iD})\} \cup \\
& \quad \quad \quad \quad \text{ct2}(t, iD, qD)) \\
& \quad)
\end{aligned}$$

Proof: This is a direct extension of Theorem 1. \blacksquare

Fig. 8 is an algorithm that uses such an iterative procedure to perform the schedulability test as expressed by Theorem 2.

V. OUR TOOL

Recall Fig. 5 presented an algorithm for performing exact schedulability analysis of ZSRM-S and Fig. 7 presented an algorithm for performing exact schedulability analysis of ZSRM-SE. These algorithms have in common that they check if a set of constraints is feasible and they also solve a problem of maximizing an objective function subject to certain constraints. Some of the constraints mentioned are not MILP — they have binary variables and logical operators. We will now discuss how to convert them to MILP. In our problems, we can add, for each real variable a the constraint: $a \leq \text{BIG}$ where BIG is a constant computed as $\text{BIG} = \sum_{\tau_i \in \tau} \lceil \frac{t}{T_i} \rceil \times C_i^o$.

Sets :

$$\begin{aligned} \text{TS} &= \{i' | (\tau_{i'} \in \tau) \wedge ((\text{prio}_{i'} \geq \text{prio}_{\text{ID}}) \vee (\zeta_{i'} \geq \zeta_{\text{ID}}))\}, \text{TSHC}(i) = \{i' | (\tau_{i'} \in \tau) \wedge (\zeta_{i'} > \zeta_i)\}, \\ \text{TSHP}(i) &= \{i' | (\tau_{i'} \in \tau) \wedge (\text{prio}_{i'} > \text{prio}_i)\}, \text{QS}(i') = \{1.. \lceil \frac{t}{T_{i'}} \rceil\}, \text{PS} = \{1..4 \times \sum_{i' \in \text{TS}} \lceil \frac{t}{T_{i'}} \rceil\} \end{aligned}$$

Constraints :

$$\begin{aligned} t^1 &= 0 & \forall p \in (\text{PS} \setminus \{\text{PS}\}) : t^p &\leq t^{p+1} : \\ & & \forall \langle i, q \rangle \text{ s.t. } (i \in \text{TS}) \wedge (q \in (\text{QS}(i) \setminus \{1\})) : A_{i,q} - A_{i,q-1} &\geq T_i \\ & & \forall \langle i, q \rangle \text{ s.t. } (i \in \text{TS}) \wedge (q \in \text{QS}(i)) : \sum_{p' \in \text{PS}} \text{arrives}_{i,q}^{p'} = 1 & \sum_{p' \in \text{PS}} \text{ZS}_{i,q}^{p'} = 1 \\ & & \forall \langle i, q, p \rangle \text{ s.t. } (i \in \text{TS}) \wedge (q \in \text{QS}(i)) \wedge (p \in (\text{PS} \setminus \{\text{PS}\})) : \sum_{p' \in \text{PS}} \text{finishes}_{i,q}^{p'} = 1 - \text{terminated}_{i,q}^p \\ (\text{arrives}_{i,q}^p = 1) &\Rightarrow (\text{donex}_{i,q}^p = 0) & (\text{arrives}_{i,q}^p = 1) \Rightarrow (A_{i,q} = t^p) & \text{arrived}_{i,q}^p = \sum_{p' \in \{1..p\}} \text{arrives}_{i,q}^{p'} \\ (\text{finishes}_{i,q}^p = 1) &\Rightarrow (\text{donex}_{i,q}^p = c_{i,q}) & (\text{finishes}_{i,q}^p = 1) \Rightarrow (f_{i,q} = t^p) & \text{fin}_{i,q}^p = \sum_{p' \in \{1..p\}} \text{finishes}_{i,q}^{p'} \\ & & (\text{ZS}_{i,q}^p = 1) \Rightarrow (A_{i,q} + Z_i = t^p) & \text{Zd}_{i,q}^p = \sum_{p' \in \{1..p\}} \text{ZS}_{i,q}^{p'} \\ (\text{doneexactly}_{i,q}^p = 1) &\Rightarrow (\text{donex}_{i,q}^p = C_i) & \sum_{p' \in \{1..p\}} \text{doneexactly}_{i,q}^{p'} = \text{highexeandnotterminated}_{i,q}^p \\ (\text{highexeandnotterminated}_{i,q}^p = 1) &\Leftrightarrow ((\text{arrived}_{i,q}^p = 1) \wedge (\text{atmostnom}_{i,q}^p = 0) \wedge (\text{terminated}_{i,q}^p = 0)) \\ & & (\text{arrivednotfin}_{i,q}^p = 1) \Leftrightarrow ((\text{arrived}_{i,q}^p = 1) \wedge (\text{fin}_{i,q}^p = 0)) \\ & & (\text{Zdnotfin}_{i,q}^p = 1) \Leftrightarrow ((\text{Zd}_{i,q}^p = 1) \wedge (\text{fin}_{i,q}^p = 0)) \\ & & (\text{atmostnom}_{i,q}^p = 1) \Leftrightarrow (\text{donex}_{i,q}^p \leq C_i) \\ & & (\text{atmostnom}_{i,q}^p = 0) \Leftrightarrow (\text{donex}_{i,q}^p > C_i) \\ (\text{Zdnotfinandnotatmostnom}_{i,q}^p = 1) &\Leftrightarrow ((\text{Zd}_{i,q}^p = 1) \wedge (\text{fin}_{i,q}^p = 0) \wedge (\text{atmostnom}_{i,q}^p = 0)) \\ & & \text{terminated}_{i,q}^p = \sum_{p' \in \{1..p\}} \text{terminatednow}_{i,q}^{p'} \\ (\text{suspendednow}_{i,q}^p = 1) &\Rightarrow (\sum_{i' \in \text{TSHC}(i)} \sum_{q' \in \text{QS}(i')} \text{Zdnotfin}_{i',q'}^p \geq 1) \\ (\text{suspendednow}_{i,q}^p = 0) &\Rightarrow (\sum_{i' \in \text{TSHC}(i)} \sum_{q' \in \text{QS}(i')} \text{Zdnotfin}_{i',q'}^p \leq 0) \\ (\text{terminatednow}_{i,q}^p = 1) &\Rightarrow (\sum_{i' \in \text{TSHC}(i)} \sum_{q' \in \text{QS}(i')} \text{Zdnotfinandnotatmostnom}_{i',q'}^p \geq 1) \\ (\text{terminatednow}_{i,q}^p = 0) &\Rightarrow (\sum_{i' \in \text{TSHC}(i)} \sum_{q' \in \text{QS}(i')} \text{Zdnotfinandnotatmostnom}_{i',q'}^p \leq 0) \\ (\text{eligZSRMSE}_{i,q}^p = 1) &\Leftrightarrow ((\text{arrivednotfin}_{i,q}^p = 1) \wedge (\text{suspendednow}_{i,q}^p = 0) \wedge (\text{terminated}_{i,q}^p = 0)) \\ (\text{candZSRMSE}_{i,q}^p = 1) &\Leftrightarrow ((\text{eligZSRMSE}_{i,q}^p = 1) \wedge (\bigwedge_{i' \in \text{TSHP}(i)} \bigwedge_{q' \in \text{QS}(i')} (\text{eligZSRMSE}_{i',q'}^p = 0))) \\ (\text{busy}^p = 1) &\Rightarrow (\sum_{\tau_{i'} \in \tau} \sum_{q' \in \text{QS}(i')} \text{candZSRMSE}_{i',q'}^p \geq 1) & (\text{busy}^p = 0) \Rightarrow (\sum_{\tau_{i'} \in \tau} \sum_{q' \in \text{QS}(i')} \text{candZSRMSE}_{i',q'}^p \leq 0) \\ & & (x_{i,q}^p = 1) \Rightarrow (\text{donex}_{i,q}^{p+1} = \text{donex}_{i,q}^p + t^{p+1} - t^p) & x_{i,q}^p \leq \text{candZSRMSE}_{i,q}^p \\ \forall p \in (\text{PS} \setminus \{\text{PS}\}) : & & \sum_{\tau_{i'} \in \tau} \sum_{q' \in \text{QS}(i')} x_{i',q'}^p = \text{busy}^p \\ \forall p \in \text{PS} : & & (\text{finishes}_{\text{ID},\text{qD}} = 1) \Rightarrow (\sum_{p' \in \{1..p-1\}} \text{busy}^{p'} \geq p - 1) \\ \forall \langle i, q \rangle \text{ s.t. } (i \in \text{TS}) \wedge (q \in \text{QS}(i)) \wedge (\zeta_i > \zeta_{\text{ID}}) : c_{i,q} \leq C_i & & \forall \langle i, q \rangle \text{ s.t. } (i \in \text{TS}) \wedge (q \in \text{QS}(i)) \wedge (\zeta_i \leq \zeta_{\text{ID}}) : c_{i,q} \leq C_i^o \\ & & (\text{terminated}_{\text{ID},\text{qD}} = 0) \Rightarrow (\text{ft} = f_{\text{ID},\text{qD}}) & \forall p \in \text{PS} : (\text{terminatednow}_{\text{ID},\text{qD}}^p = 1) \Rightarrow (\text{ft} = t^p) \\ \text{Domains of variables} : t^p \in \mathbb{R}_{\geq 0}, A_{i,q}^p \in \mathbb{R}_{\geq 0}, \text{arrives}_{i,q}^p \in \{0, 1\}, \text{arrived}_{i,q}^p \in \{0, 1\}, f_{i,q} \in \mathbb{R}_{\geq 0}, \text{finishes}_{i,q}^p \in \{0, 1\}, \\ \text{fin}_{i,q}^p \in \{0, 1\}, \text{ZS}_{i,q}^p \in \{0, 1\}, \text{Zd}_{i,q}^p \in \{0, 1\}, \text{donex}_{i,q}^p \in \mathbb{R}_{\geq 0}, c_{i,q} \in \{0, 1\}, \text{arrivednotfin}_{i,q}^p \in \{0, 1\}, \text{Zdnotfin}_{i,q}^p \in \{0, 1\}, \\ \text{suspendednow}_{i,q}^p \in \{0, 1\}, \text{eligZSRMSE}_{i,q}^p \in \{0, 1\}, \text{candZSRMSE}_{i,q}^p \in \{0, 1\}, \text{busy}^p \in \{0, 1\}, x_{i,q}^p \in \{0, 1\}, \text{ft} \in \mathbb{R}_{\geq 0} \end{aligned}$$

Fig. 7: Constraints we use for exact schedulability analysis of ZSRM-SE.

```

1. allOK := true
2. for each  $\tau_{iD} \in \tau$ , as long as allOK do
3.    $t := -1$ ;  $newt := C_{iD}^o$ 
4.   while ( $t < newt$ ) and allOK do
5.      $t := newt$ 
6.     for each  $qD \in \{1.. \lceil \frac{t}{\tau_{iD}} \rceil\}$ , as long as allOK do
7.       if  $feas(\{(terminated_{iD,qD}^{PS} = 1) \vee$ 
8.          $(f_{iD,qD} - A_{iD,qD} > D_{iD})\} \cup ct2(t, iD, qD))$  then
9.         allOK := false
10.      end if
11.    end for
12.    if allOK then
13.      flag := false
14.      for each  $qD' \in \{1.. \lceil \frac{t}{\tau_{iD}} \rceil\}$  do
15.         $\langle fe, va \rangle :=$ 
16.          solve(max{ft|ct2(t, iD, qD')})
17.        if fe then
18.          if flag then  $newt := \max(newt, va)$ 
19.          else  $newt := va$ ; flag := true
20.          end if
21.        end if
22.      end for
23.    end while
24.  end for
  return allOK

```

Fig. 8: An algorithm for ZSRM-SE schedulability testing; it is optimized for detecting clearly unschedulable tasksets quickly.

And this does not change feasibility. A constraint of the form $(x = 1) \Rightarrow (a = b)$ can be rewritten as: $((x = 1) \Rightarrow (a \leq b)) \wedge ((x = 1) \Rightarrow (a \geq b))$. Note that if x is a variable with the domain $\{0, 1\}$ and a and b are non-negative real variables and BIG is a constant selected so that $a \leq BIG$ and $b \leq BIG$, then a constraint $(x = 1) \Rightarrow (a \leq b)$ can be rewritten as

$$a - b + BIG \times x \leq BIG \quad (5)$$

Also, a constraint of the form $(a = 1) \Leftrightarrow (b = 1) \wedge (c = 1)$ can be rewritten as $(b + c - a \leq 1) \wedge (b + c - 2a \geq 0)$. Consider the constraint: $(suspendednowZSRMS_{i,q}^p = 1) \Rightarrow (\sum_{i' \in TSHC(i)} \sum_{q' \in QS(i')} Z_{\text{notfin}}ZSRMS_{i',q'}^p \geq 1)$. When we rewrite it, we use $BIG = 1 + (\sum_{i' \in TSHC(i)} |QS(i')|)$.

With these techniques, we can rewrite the optimization problems and feasibility checking problems are MILP. Indeed, we have done so and implemented a tool that performs these computations. Our implementation uses Gurobi 6.0.3 — a state-of-the-art MILP solver.

VI. CONCLUSIONS

Zero-Slack Rate-Monotonic (ZSRM) is a mixed-criticality scheduler which suspends a low-criticality tasks when a high-criticality tasks has not finished at a certain time. Previous work has made available implementations of a ZSRM scheduler in the Linux kernel and in VxWorks, as well as a sufficient schedulability test for it and this schedulability test is available to software practitioners in the OSATE AADL workbench. And a modification of it was used in a UAV system to ensure that an overload in vision processing does not jeopardize deadline guarantees of flight-control software [6]. Hence, we

believe ZSRM is one of the most practical ideas in mixed-criticality scheduling. Unfortunately, no exact schedulability analysis was available for ZSRM schedulers. Therefore, in this paper, we presented exact schedulability tests for two ZSRM schedulers.

ACKNOWLEDGMENT

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. This material has been approved for public release and unlimited distribution. DM-0002431

REFERENCES

- [1] S. K. Baruah, A. Burns, and R. I. Davis. Response-time analysis for mixed criticality systems. In *RTSS'11*.
- [2] S. K. Baruah and S. Vestal. Schedulability analysis of sporadic tasks with multiple criticality specifications. In *ECRTS'08*.
- [3] A. Burns and R. I. Davis. Mixed criticality systems — a review. Technical report, Department of Computer Science, University of York, 2015. www-users.cs.york.ac.uk/burns/review.pdf.
- [4] R. I. Davis and A. J. Wellings. Dual priority scheduling. In *RTSS'95*.
- [5] D. de Niz, K. Lakshmanan, and R. Rajkumar. On the scheduling of mixed-criticality real-time task sets. In *RTSS'09*.
- [6] D. de Niz, L. Wrage, A. Rowe, and R. Rajkumar. Utility-based resource overbooking for cyber-physical systems. In *RTCSA'13*.
- [7] H.-M. Huang, C. Gill, and C. Lu. Implementation and evaluation of mixed criticality scheduling approaches for periodic tasks. In *RTAS'12*.
- [8] J. Lehoczky, L. Sha, and Y. Ding. The rate monotonic scheduling algorithm: Exact characterization and average case behavior. In *RTSS'90*.
- [9] J. P. Lehoczky, L. Sha, and J. K. Strosnider. Enhanced aperiodic responsiveness in hard real-time environments. In *RTSS'87*.
- [10] C. L. Liu and J. W. Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *J. ACM*, 20(1):46–61, 1973.
- [11] L. Sha, J. P. Lehoczky, and R. Rajkumar. Solutions for some practical problems in prioritized preemptive scheduling. In *RTSS'86*.
- [12] S. Vestal. Preemptive scheduling of multi-criticality systems with varying degrees of execution time assurance. In *RTSS'07*.