

# Mixed Criticality Systems - A Review\*

Alan Burns and Robert I. Davis  
Department of Computer Science,  
University of York, York, UK.  
email: {alan.burns, rob.davis}@york.ac.uk

## Abstract

This review covers research on the topic of mixed criticality systems that has been published since Vestal's 2007 paper. It covers the period up to and including July 2018. The review is organised into the following topics: introduction and motivation, models, single processor analysis (including job-based, hard and soft tasks, fixed priority and EDF scheduling, shared resources and static and synchronous scheduling), multiprocessor analysis, related topics, realistic models, formal treatments, systems issues and industrial practice.

---

\*eleventh edition, August 2018.

# Contents

<b>1</b>	<b>Introduction</b>	<b>4</b>
<b>2</b>	<b>Mixed Criticality Models</b>	<b>6</b>
<b>3</b>	<b>Single Processor Analysis</b>	<b>8</b>
3.1	Job scheduling . . . . .	8
3.2	Fixed Priority Scheduling . . . . .	9
3.2.1	RTA-Based approaches . . . . .	9
3.2.2	Slack scheduling . . . . .	13
3.2.3	Period transformation . . . . .	14
3.3	EDF Scheduling . . . . .	15
3.4	Shared Resources . . . . .	17
3.5	Static and Synchronous Scheduling . . . . .	18
3.6	Varying Speed Processors . . . . .	19
<b>4</b>	<b>Multiprocessor Analysis</b>	<b>19</b>
4.1	Task Allocation . . . . .	20
4.2	Schedulability Analysis . . . . .	22
4.3	Communication and other Resources . . . . .	22
<b>5</b>	<b>Links to other Research Topics</b>	<b>26</b>
5.1	Hard and Soft Tasks . . . . .	26
5.2	Fault Tolerant Systems (FTS) . . . . .	28
5.3	Security . . . . .	29
5.4	Hierarchical Scheduling . . . . .	29
5.5	Cyber Physical Systems and Internet of Things . . . . .	30
5.6	Probabilistic real-time systems . . . . .	31
<b>6</b>	<b>More Realistic MCS Models</b>	<b>33</b>
<b>7</b>	<b>More Formal Treatments</b>	<b>37</b>
7.1	Utilisation Bounds . . . . .	37
7.2	Speedup Factors . . . . .	37
7.3	Formal Language and Modelling Issues . . . . .	38
<b>8</b>	<b>Systems Issues</b>	<b>39</b>
<b>9</b>	<b>Industry Practice and Standards</b>	<b>44</b>

<b>10 PhD Awards for Research within the Field of MCS</b>	<b>45</b>
<b>11 Conclusion and Directions for Future Work</b>	<b>46</b>

# 1 Introduction

An increasingly important trend in the design of real-time and embedded systems is the integration of components with different levels of criticality onto a common hardware platform. At the same time, these platforms are migrating from single cores to multi-cores and in the future many-core architectures. Criticality is a designation of the level of assurance against failure needed for a system component. A mixed criticality system (MCS) is one that has two or more distinct levels (for example safety critical, mission critical and low-critical). Perhaps up to five levels may be identified (see, for example, the IEC 61508, DO-178B and DO-178C, DO-254 and ISO 26262 standards). Typical names for the levels are ASILs (Automotive Safety and Integrity Levels), DALs (Design Assurance Levels or Development Assurance Levels) and SILs (Safety Integrity Levels). It should be noted that not all standards and papers on MCS assign the same meaning to ‘criticality’, an issue explored by Graydon and Bate [204], Esper et al. [173], Paulitsch et al. [374] and Ernst and Di Natale [172].

Most of the complex embedded systems found in, for example, the automotive and avionics industries are evolving into mixed criticality systems in order to meet stringent non-functional requirements relating to cost, space, weight, heat generation and power consumption (the latter being of particular relevance to mobile systems). Indeed the software standards in the European automotive industry (AUTOSAR<sup>1</sup>) and in the avionics domain (ARINC<sup>2</sup>) address mixed criticality issues; in the sense that they recognise that MCS must be supported on their platforms.

The fundamental research question underlying these initiatives and standards is: how, in a disciplined way, to reconcile the conflicting requirements of *partitioning* for (safety) assurance and *sharing* for efficient resource usage. This question gives rise to theoretical problems in modelling and verification, and systems problems relating to the design and implementation of the necessary hardware and software run-time controls.

A key aspect of MCS is that system parameters, such as tasks’ worst-case execution times (WCETs), become dependent on the criticality level of the tasks. So the same code will have a higher WCET if it is defined to be safety-critical (as a higher level of assurance is required) than it would if it is just considered to be mission critical or indeed non-critical. This property of MCS significantly modifies/undermines many of the standard scheduling results. This report aims to review the research that has been published on MCS.

The first paper on the verification of a Mixed Criticality System used an exten-

---

<sup>1</sup><http://www.autosar.org/>

<sup>2</sup><http://www.arinc.com/>

sion of standard fixed priority (FP) real-time scheduling theory, and was published by Vestal (of Honeywell Aerospace) in 2007 [451]<sup>3</sup>. It employed a somewhat restrictive work-flow model, focused on a single processor and made use of Response Time Analysis [28]. It showed that neither rate monotonic [319] nor deadline monotonic [304] priority assignment is optimal for MCS; however Audsley's optimal priority assignment algorithm [26] was found to be applicable.

This paper was followed by two publications in 2008 by Baruah and Vestal [75], and Huber et al. [255]. The first of these papers generalises Vestal's model by using a sporadic task model and by assessing fixed job-priority scheduling and dynamic priority scheduling. It contains the important result that EDF (Earliest Deadline First) does not dominate FP when criticality levels are introduced, and that there are feasible systems that cannot be scheduled by EDF. The latter paper addresses multi-processor issues and virtualisation (though it did not use that term). It focused on AUTOSAR and resource management (encapsulation and monitoring) with time-triggered applications and a trusted network layer.

Further impetus to defining MCS as a distinct research topic came from the white paper produced by Barhorst et al. [40], the keynote talk that Baruah gave at the 2010 ECRTS conference<sup>4</sup> and a workshop report from the European Commission [441]. These have been followed up by tutorials on MCS at ESWEEK in 2012 and 2013<sup>5</sup>, a workshop at HiPEAC in January 2013<sup>6</sup>, a workshop (WICERT) at DATE 2013<sup>7</sup>, a workshop (ReTiMiCS) at RTCSA 2013<sup>8</sup>, workshops (WMC) at RTSS 2013<sup>9</sup>, RTSS 2014<sup>10</sup>, RTSS 2015<sup>11</sup>, RTSS 2016<sup>12</sup> and RTSS 2017<sup>13</sup>; a workshop at the 19th International Conference on Reliable Software Technologies (Ada-Europe) in June 2014, and Dagstuhl Seminars on Mixed Criticality and Many

---

<sup>3</sup>The term Mixed Criticality had been used before 2007 to address issues of non-interference in non-federated architectures such as IMA [239]; Vestal changed the focus of research by concentrating on real-time performance. Systems with more than one criticality level but aim to only give complete isolation are called *multiple-criticality systems*; the use of *mixed-criticality* implies some tradeoff between isolation and integration that involves resource sharing.

<sup>4</sup>Available from the conference web site: <http://ecrts.eit.uni-kl.de/index.php?id=53>.

<sup>5</sup>Embedded Systems Week: <http://www.esweek.org/>

<sup>6</sup><http://www.hipeac.net/conference/berlin/workshop/integration-mixed-criticality-subsystems-multi-core-processors>

<sup>7</sup><http://atcproyectos.ugr.es/wicert/index.php/conference-proceedings>

<sup>8</sup><http://igm.univ-mlv.fr/rtalgo/Events/RETIMICS/>

<sup>9</sup><http://www.cs.york.ac.uk/robdavis/wmc2013/>

<sup>10</sup><http://www.cs.york.ac.uk/robdavis/wmc2014/>

<sup>11</sup><http://www.cs.york.ac.uk/robdavis/wmc/>

<sup>12</sup><https://gsathish.github.io/wmc2016/>

<sup>13</sup><https://cps-research-group.github.io/WMC2017/>

Core Platforms in 2015<sup>14</sup> and 2017<sup>15</sup>.

This review [103] is organised as follows. In Section 2 we first consider mixed criticality models. Then in Section 3 single processor systems are covered (including fixed priority and EDF scheduling). Section 4 covers multiprocessor issues and Section 5 links this research to other topics such as hard and soft real-time scheduling and hierarchical scheduling. More realistic models are covered in Section 6, more formal work is covered in Section 7 and systems work is covered in Section 8. Industry practice and safety standards provide a somewhat different perspective on MCS to Vestal’s model; these differences are discussed in Section 9. Section 10 lists PhD dissertations that have been produced on MCS since 2014. The review concludes with Section 11 which outlines a number of open problems and areas where further research is needed.

An adaptation of this review, covering publications up to the end of 2016, has been published in ACM Computer Surveys [105]. This should be used as the main citation for this report. Other overviews/surveys on MCS have also been produced [21, 220].

## 2 Mixed Criticality Models

Inevitably not all papers on mixed criticality have used the same system or task model. Here we define a model that is generally applicable and is capable of describing the main results considered in this review.

A system is defined as a finite set of components  $\mathcal{K}$ . Each component has a level of criticality (designated by the systems engineer responsible for the entire system),  $L$ , and contains a finite set of *sporadic tasks*. Each task,  $\tau_i$ , is defined by its period (minimum arrival interval), deadline, computation time and criticality level:  $(T_i, D_i, C_i, L_i)$ . Tasks give rise to a potentially unbounded sequence of *jobs*.

The primary concern with the implementation of MCS is one of separation. Tasks from different components must not be allowed to interfere with each other. In particular, mechanisms must be in place to prevent a job from executing for more than the computation time  $C$  defined for its task, and to ensure that a task does not generate jobs that are closer together than  $T$ <sup>16</sup>.

The requirement to protect the operation of one component from the faults of another is present in all systems that host multiple applications. It is however of particular significance if components have different criticality levels. Since without

---

<sup>14</sup><http://www.dagstuhl.de/15121>

<sup>15</sup><http://www.dagstuhl.de/17131>

<sup>16</sup>Or (period minus release jitter) if that is part of the task model.

such protection, all components would need to be engineered to the strict standards of the highest criticality level, potentially massively increasing development costs.

After concerns of partitioning comes the need to use resources efficiently. This is facilitated by noting that the task parameters are not independent, in particular the worst-case computation/execution time estimate,  $C_i$ , will be derived by a process dictated by the criticality level. The higher the criticality level, the more conservative the verification process and hence the greater will be the value of  $C_i$ . This was the observation at the heart of the paper by Vestal [451].

For systems executing on hardware platforms with deterministic behaviour, any particular task will have a single real WCET (worst-case execution time); however, this value typically cannot be known with complete certainty. This uncertainty is primarily epistemic (uncertainty in what we know, or do not know, about the system) rather than aleatory (uncertainty in the system itself). Although it is reasonable to assume confidence increases (i.e. uncertainty decreases) with larger estimates of worst-case execution time, this may not be universally true [204]. It would certainly be hard to estimate what increase in confidence would result from, say, a 10% increase in all  $C$ s.

For systems executing on hardware platforms with time-randomised hardware components [115], then a probabilistic WCET (pWCET) [18, 137, 147, 164] can be obtained. The exceedance function for this probability distribution defines for any specific probability, derived from a required maximum failure rate associated with a criticality level, an execution time budget which has no greater probability of being exceeded on any given run [136]. The pWCET distribution therefore effectively defines different estimates of the WCET budget for the same task, for different criticality levels due to their different requirements on the maximum tolerable failure rate.

The focus on different computation times was extended to task periods in subsequent papers [44, 48, 50, 58, 63, 97, 101, 472]. Here tasks are event handlers. The higher the criticality level the more events must be handled, and hence the task must execute more frequently even if it does not execute for longer.

In MCS a task is now defined by:  $(\vec{T}, D, \vec{C}, L)$ , where  $\vec{C}$  and  $\vec{T}$  are vectors of values – one per criticality level, with the constraints:

$$L1 > L2 \Rightarrow C(L1) \geq C(L2)$$

$$L1 > L2 \Rightarrow T(L1) \leq T(L2)$$

for any two criticality levels  $L1$  and  $L2$ .

Note the completion of the model, by making  $D$  criticality dependent [58] has not as yet been addressed in detail. But it could have the constraint:

$$L1 > L2 \Rightarrow D(L1) \geq D(L2)$$

So a task may have a ‘safety critical’ deadline and an early Quality of Service (QoS) deadline. Alternatively:

$$L1 > L2 \Rightarrow D(L1) \leq D(L2)$$

in which case the conservative ‘safety critical’ deadline is shorter than the one deemed necessary if the criticality level is lower.

Another feature of many of the papers considered in this review is that the system is defined to execute in a number of criticality *modes*. A system starts in the lowest criticality mode. If all jobs behave according to this mode then the system stays in that mode. But if any job attempts to execute for a longer time, or more frequently, than is acceptable in that mode then a criticality mode change occurs. Ultimately the system may change to the highest criticality mode.

Some papers allow the criticality mode to move down as well as up, but others (indeed the majority) restrict the model to increases in criticality only. We return to this issue in Section 6.

Finally, many papers restrict themselves to just two criticality levels; high (*HI*) and low (*LO*) with  $HI > LO$ . These are referred to as *dual-criticality* systems. Where modes are used, the system is either in a *LO*-criticality (or normal) mode or a *HI*-criticality mode. And the set of task parameters is typically:  $(T_i, D_i, C_i(HI), C_i(LO), L_i)$ . At the other extreme are the models presented by Ekberg et al. [165, 167, 169] in which any number of modes are allowed and the movement between modes is represented by a directed acyclic graph.

### 3 Single Processor Analysis

Since Vestal’s 2007 paper [451] there has been a series of publications. Most of these papers address single processor platforms and independent components.

#### 3.1 Job scheduling

Initially a number of papers considered the restricted problem of scheduling, on a single processor, a finite set of mixed criticality jobs with criticality dependent execution times [43, 51, 53, 67, 69, 74, 79, 210, 306, 307, 369, 408, 416, 417, 421]. This work has, however, largely been superseded by work on the more widely applicable task model.



## 3.2 Fixed Priority Scheduling

In this section we look at MCS schemes that are based on applying Response-Time Analysis (RTA), then those that consider slack scheduling and finally approaches that are derived from period transformations.

### 3.2.1 RTA-Based approaches

Vestal's approach was formalised (i.e. proof that the use of Audsley's priority assignment algorithm [26] was optimal) by Dorin et al. [156] in 2010. They also extended the model to include release jitter, and showed how sensitivity analysis could be applied.

Vestal's approach allowed the priorities of high and low criticality tasks to be interleaved, but all tasks had to be evaluated as if they were of the highest criticality. By introducing monitoring of task execution time, and the prevention of execution time over-runs, higher resource usage can be delivered [58]. This is a crucial issue in mixed criticality scheduling; by the introduction of more trusted components a high utilisation of the available resources is facilitated.

In 2011 this approach was further extended [61,97] to give a scheduling model and associated analysis framework for a single processor system that dominates all previous published analysis for MCS (using fixed priority scheduling) in that it made better use of the processor and could schedule all systems that could be guaranteed by other approaches, plus many that could not. These papers were however restricted to just two criticality levels (or modes). The system's run-time behaviour is either *low-criticality* (which relies on all execution times being bound by the low-criticality values and guarantees that all deadlines are met) or *high-criticality* (where only high criticality work is guaranteed but the rely condition<sup>17</sup> is weakened – the bound on high-criticality execution times is increased). The system's criticality change (from Low to High, i.e. *LO* to *HI*) is triggered by the observation, at run-time, that the stronger rely condition has been violated. In the context of control applications Cheng et al [120] also allow the state of the controlled plant to trigger a criticality mode change.

This change in criticality level has a number of similarities to systems that move between different operational modes (although there are also some significant differences [95, 204]). In the *HI*-criticality mode there are fewer tasks, but they have longer execution times or shorter periods. The literature on mode change protocols [37, 108, 170, 375, 390, 412, 442, 443], however, highlights one important problem: a system can be schedulable in every mode, but not schedulable during a mode change [443]. This is also true for systems that change criticality levels.

---

<sup>17</sup>A rely condition formalises the assumptions required for the guarantees to be valid [268].

An optimal priority ordering is defined in the paper from Baruah et al. [61] in that it maximises the priority of high criticality tasks, subject to the system being schedulable. Both the high and low criticality tasks are ordered via deadline (deadline monotonic) and a simplified version of Audsley’s algorithm is used to assign priorities from the lowest to the highest level. At each priority level the lowest priority task from the low criticality task set is tried first, if it is schedulable then the algorithm moves up to the next priority level; if it is not schedulable then the lowest priority task from the high criticality set is tested. If it is schedulable then again the algorithm moves on to the next level. But if neither of these two tasks are schedulable then the search can be abandoned as the task set is unschedulable. In total a maximum of  $2N$  tests are needed (where  $N$  is the number of tasks in the system)<sup>18</sup>. Note that this result follows from work on robust priority assignment [145]. As each set of LO/HI criticality tasks can be viewed as additional interference on the other subset, an optimal priority ordering can be obtained with each subset in Deadline Monotonic priority order and a merge operation between them.

The protocol (dropping all *LO*-criticality work if any task executes for more than its  $C(LO)$  value<sup>19</sup>), the derived analysis and the use of optimal priority ordering is shown [61] to out-perform other schemes (in terms of success ratio for randomly generated task sets). The analysis is based on standard RTA (Response-Time Analysis). For any task,  $\tau_i$ , first its *LO*-criticality response-time ( $R(LO)$ ) is computed using *LO*-criticality parameters for all the tasks. A criticality switch must occur before this value if the task is to be impacted by the change, otherwise it will have completed execution. The worst-case response-time in the *HI*-criticality mode ( $R(HI)$ ) is computed by noting that all *LO*-criticality tasks must be abandoned by time  $R(LO)$ . The paper contains two methods for computing  $R(HI)$ <sup>20</sup>, one involves a single upper bound, the other looks at all the possible critically change points before  $R(LO)$  and computes the worst-case. The latter is more accurate, though still not exact; however, the gain in performance is not significant and the simple upper bound test is probably sufficient in most cases.

To illustrate the above approaches one of the graphs from [61] is reproduced in Figure 1. This figure plots the percentage of task sets generated that were deemed schedulable for a system of 20 tasks, with on average 50% of those tasks having high criticality and each task having a high criticality execution time that is twice its low criticality execution time. The compared approaches are (from least effective to most effective): CrMPO which assigned priorities in criticality order, SMC-NO

<sup>18</sup>Strictly, only  $2N-1$  tests are needed as the highest priority task must be schedulable as its computation time is less than its deadline.

<sup>19</sup>First proposed by Baruah [43, 51].

<sup>20</sup>In [61] the example in Section IV.B (final step) should have a worst-case response-time of 90, not 85 as reported in the paper; however 90 is still below the deadline of 100.




Figure 1: Percentage of Schedulable Task Sets

(static mixed criticality with no run-time monitoring) which is Vestal’s original approach, SMC which is an adaptation of Vestal’s approach in which  $LO$ -criticality tasks are monitored at run-time and are prevented from executing for more than  $C(LO)$ , and AMC-rtb and AMC-max which are the two methods introduced in the previous paragraph (AMC for adaptive mixed criticality). In the graph the UB-H&L line bounds the maximum possible number of schedulable task sets. It serves to illustrate the quality of the AMC-max approach. Almost all publications on MCS assume that the task deadlines are constrained ( $D \leq T$ ). For AMC, Burns and David [104] removed this constraint and a version of the analysis for AMC now allows arbitrary deadlines.

The AMC-rtb approach was extended by Zhao et al [475, 476, 479] in 2013 to incorporate preemption thresholds [398] into the model. They demonstrated a reduction in stack usage and improved performance for some parameter ranges. Another approach to combining AMC-rtb and existing scheduling theory is taken by Burns and Davis [102]. They consider the use of deferred preemption [93, 142] and demonstrate a significant improvement over fully preemptive AMC-rtb. The gain in schedulability they demonstrate is obtained by having a final non-preemptive region (FNPR) at the end of  $C(LO)$  and  $C(HI)$ , and by combining the assignment of

priority and the determination of the size of these FNPRs. Where scheduling analysis is part of a design optimisation Zhao and Zeng [480] argue that even AMC-rtb is too complex. They propose a new simpler test that is still safe and has bounded pessimism. At the other extreme, Asyaban and Kargahi [23] develop exact analysis for AMC, unfortunately at the cost of losing optimal priority ordering.

In keeping with a number of papers on MCS, the work of Baruah et al. [61] (and most of the subsequent modifications) restricted itself to dual criticality systems. Fleming and Burns [182] extended these models to an arbitrary number of criticality levels, focusing particularly on five levels as this is the maximum found in automotive and avionics standards. They observed that AMC-rtb remains a good approximation to AMC-max, and that AMC-max became computationally expensive for increased numbers of criticality levels. They concluded that AMC-rtb represented an adequate and effective form of analysis. A relatively minor improvement to AMC-max was published by Huang et al. [246] (they termed it AMC-IA); however there are cases where their analysis is optimistic (i.e. unsound) [186].

One characteristic of all the schemes defined above is that tasks do not change their priority after a criticality mode change. If priorities can change then a simple form of sufficient analysis is possible [62]. This work defines a new approach, PMC (priority may change). Evaluations show that PMC performs similarly to AMC-rtb, though neither dominates the other. An improved scheme, GFP (Generalised Fixed Priority) is proposed by Chen et al. [119]. They assign (using an heuristic) three priorities to each task. One for each of the two criticality levels, and one for the transition between the criticality modes. They demonstrate an improvement over AMC-rtb.

It was noted in the section on Mixed Criticality Models that the period parameter ( $T$ ) can be criticality dependent as well as the worst-case execution time estimate ( $C$ ). An application may consist of event handlers, and have different levels of constraint over the arrival patterns of the events. The higher the criticality, the closer together the events are assumed to arrive; and hence the smaller the  $T$  parameter. Baruah and Chattopadhyay [63] have reformulated the SMC and AMC analysis (introduced above) to apply to this model, in which the  $T$ s rather than the  $C$ s vary with criticality. Their evaluation results show similar behaviour to that depicted in Figure 1. Criticality specific periods are also addressed by Burns and Davis [101], Baruah [50], and by Zhang et al. [472] (who derived an improved analysis that they termed SAMC – Sufficient AMC).

For periodic task sets with offsets, Asyaban et al. [24] has produced feasibility analysis showing that simulation over an interval of length four times the hyperperiod plus the largest offset forms a sufficient test of schedulability.

### 3.2.2 Slack scheduling

An alternative approach to scheduling mixed criticality fixed priority systems is, for dual-criticality systems, to use a *slack scheduling* scheme in which low criticality jobs are run in the slack generated by high criticality jobs only using their low criticality execution budgets. This was first explored by Niz et al. [356]. One difficulty with this approach is to incorporate sporadic tasks. At what point can the ‘slack’ of a non-appearing sporadic task be allocated to low criticality jobs? Even for periodic tasks, ensuring schedulability of high criticality tasks in all circumstances is not straightforward. Niz et al. [356] compute the time at which a high criticality task must be released to ensure that it meets its deadline (a scheme similar to the dual-priority approach outlined in Section 5.1). However, Huang et al. [245] demonstrated that if a low criticality (high priority) task executes beyond its deadline, a high criticality (lower priority) task could miss its deadline. They show that either the low criticality task must be aborted at its deadline or (more practically) its priority must be reduced to a background level. They then derive safe analysis. Niz et al. subsequently modified the enforcement rule in their model to remove the problem and improved its performance [357, 358].

While slack is usually generated by tasks not executing for their full budget, it is also produced by the arrival of jobs being less frequent than anticipated in the worst-case. Neukirchner et al. [351, 352] adapt and extend a number of schemes for monitoring activation patterns. Their multi-mode approach is proved to be safe (no false negatives) and efficient (few false positives). Hu et al. [244] also consider budget management, and produce an effective scheme for minimising the overheads associated with slack management.

For a dual-criticality system  $C(LO)$  values must, of course, be known. Once schedulability has been established however, it is possible to derive [404], using sensitivity analysis [84, 383], a scaling factor  $F$  ( $F > 1$ ) such that the system remains schedulable with all  $C(LO)$  values replaced by  $F \cdot C(LO)$ . Using these scaled values at run-time will increase the robustness of the system, as the  $LO$ -criticality tasks will be able to execute for a greater time before a criticality change is triggered. Scaling can also be applied to the  $C(HI)$  values. Volp et al. [456] look at an alternative means of obtaining  $C(LO)$  and  $C(HI)$  values; they do not consider them to be estimates of worst-case execution time, but budgets set by some design optimisation process.

As scaling involves changing a task’s computation time, and computation time influences priority assignment, it is possible to extend this approach by also allowing priorities to change as the system is made more robust [98]. A more dynamic budget management scheme is used by Gu and Easwaran [212] to postpone criticality level mode changes within the context of the EDF-VD scheme (see Section 3.3).

Sciandra et al. [407] are extending and applying scaling factors to intelligent transport systems. Issues of robustness are also addressed by Herman et al. [235].

### 3.2.3 Period transformation

As Vestal noted [451], an older protocol *period transformation* [410, 411] (PT), is also applicable to the mixed criticality scheduling problem. Period transformation splits a task with period  $T$  and computation time  $C$  into two (or more) parts so that the task now has the parameters  $T/2$  and  $C/2$ . Assuming all tasks have deadlines equal to their periods, the application of the optimal rate monotonic priority assignment scheme [319] will increase the relative priority of all transformed tasks. If all high criticality tasks are transformed so that their transformed periods are shorter than all low criticality tasks then the rate monotonic algorithm will deliver partitioned (i.e. criticality monotonic) priorities. All high criticality tasks will have priorities greater than all lower criticality tasks. The scheme can easily be extended to task sets with constrained deadlines ( $D < T$ ). However, the scheme does introduce extra overheads from the increased number of context switches, and these could be excessive if there are low criticality tasks with short deadlines. A simple example of a period transformed task would be one with  $T = D = 16$ ,  $C(HI) = 8$  and  $C(LO) = 4$ ; this task could be transformed to one with  $T = D = 4$  and  $C = 2$ . Note, this is  $C(HI)/4$ , not  $C(LO)/4$ . The computation time is such that if the task executes according to its *HI*-criticality parameter it will take four invocations of the transformed task to complete, but if the *LO*-criticality assumption is valid it will only take two.

If overheads are ignored then Period Transformation performs well. Baruah and Burns postulate [59] (and prove for two tasks) that this is primarily due to the inherent property of PT to deliver tasks sets with harmonic periods (that are then more likely to be schedulable). It does not seem that PT is of specific benefit to MCS.

To split the code of a task, either a static code transformation process must be used or the run-time must employ an execution-time server. With code transformation, the programmer must identify where in the code the split should be made. This does not lead to good code modularisation and is similar to the problems encountered when functions must to be split into short sections so that they can be ‘packed’ into the minor cycles of a cyclic executive [109]. There is also the problem of OS locks being retained between slices of the code; making the protected resource unavailable to other tasks.

With a dual-criticality task such as the one in the example above the point at which the task can be assumed to have executed for two units of time is itself criticality dependent. This to all intents and purposes makes code transformation im-

practical. Therefore, if the code is not to be changed then a run-time server must be used to restrict the amount of computation allowed per release of the (transformed) task. In practice this means that:

- Without PT, *LO*-criticality tasks may have high priorities and hence their execution times must be monitored (and enforced); *HI*-criticality tasks must also be monitored as they may need to trigger a criticality change if they execute for more than  $C(LO)$  thereby triggering the abandonment of *LO*-criticality tasks.
- With PT, *LO*-criticality tasks have the lower priorities and hence they do not need to be monitored, *HI*-criticality tasks must be monitored to enforce the per release budget.

In general, there is less run-time intervention with PT. But recall there is considerably more task switching overhead if the periods of all *HI*-criticality tasks are reduced to less than all *LO*-criticality task periods.

For multiple criticality levels a number of transformations may be required to generate a criticality monotonic ordering [182]. For example if there are three tasks (H, M, and L) with criticality levels implied by their names, and periods 5, 33 and 9. Then first M must be divided by 11 to get a period of 3 (so less than 9), but then H must be divided by 5 to move it below the new value for M. As a result the transformed periods become 1, 3 and 9. It also seems that the theoretical benefit of PT diminishes with an increased number of criticality levels [182].

### 3.3 EDF Scheduling

The first paper to consider MCS with EDF scheduling was Baruah and Vestal [75] in 2008. Park and Kim [369] later introduced a slack-based mixed criticality scheme for EDF scheduled jobs which they called CBEDF (Criticality Based EDF). In essence they use a combination of off- and on-line analysis to run *HI*-criticality jobs as late as possible, and *LO*-criticality jobs in the generated slack. In effect they are utilising an older protocol developed by Chetto and Chetto [121] for running soft real-time tasks in the ‘gaps’ produced by running hard real-time tasks so as to just meet their deadlines.

A more complete analysis for EDF scheduled systems was presented by Ekberg and Yi [166, 216]. They mimicked the FP scheme by assigning two relative deadlines to each high criticality task. One deadline is the defining ‘real’ deadline of the task, the other is an artificial earlier deadline that is used to increase the likelihood of high criticality tasks executing before low criticality ones. At the point that the criticality of the system changes from low to high (due to a task exceeding its low

criticality budget), all low criticality tasks are abandoned and the high criticality tasks revert to their defining deadlines. They demonstrate a clear improvement over previous schemes [217]. Later work [167] generalises the model to include changes to all task parameters and to incorporate more than two criticality levels. Tighter analysis is provided by Easwaran [161], although it is not clear that the method will scale to more than two criticality levels. Further improvements are presented by Yao et al. [465]. They use an improved schedulability test for EDF (a scheme called QPA [471]), and a genetic algorithm (GA) to find better artificial deadlines.

A similar scheme was presented by Baruah et al. [54,56], called EDF-VD (EDF - with virtual deadlines). Again for a dual-criticality system, *HI*-criticality tasks have their deadlines reduced (if necessary) during *LO*-criticality mode execution. All deadlines are reduced by the same factor. They demonstrate both theoretically and via evaluations that this is an effective scheme. Note, however, that this scheme is not as general as those reported above [161, 166, 216]. In these approaches a different reduction factor is used for each task. Nevertheless the use of a single value does allow schedulability bounds to be derived (see Section 7). An intermediate approach that uses just two scaling factors is provided by Masrur et al. [330]; there motivation being to develop an efficient scheme that could be used at run-time. In later work [76] Baruah has generalised the underlying MCS model to include criticality-specific values for period and deadline as well as WCET. EDF-VD was further improved by Gu and Easwaran [213] by the development of a new scheduling test.

EDF scheduling of MCS is also addressed by Lipari and Buttazzo [318] using a reservation-based approach. Here sufficient budget is reserved for the high criticality tasks, but if they only make use of what is assumed by their low criticality requirements then a set of low criticality tasks can be guaranteed. Again only two criticality levels are assumed. In effect low criticality tasks run in capacity reclaimed from high criticality tasks. Deadlines for the high criticality tasks are chosen to maximise the amount of capacity reclaiming.

A different approach to using spare capacity was derived by Su et al. [429,430] by exploiting the *elastic task model* [111] in which the period of a task can change. They propose a minimum level of service for each *LO*-criticality task  $\tau_i$  that is defined by a maximum period,  $T_i^{max}$ . The complete system must be schedulable when all *HI*-criticality tasks use their  $C(HI)$  values and all *LO*-criticality tasks use their  $C(LO)$  and  $T^{max}$  values. At run-time if *HI*-criticality tasks use less than their full *HI*-criticality entitlement then the *LO*-criticality tasks can run more frequently. They demonstrate that for certain parameter sets their approach performs better than EDF-VD.

Alternative analysis for EDF scheduled MCS is presented by Santinelli et al. [400].



Their use of multiple demand-bound curves allows sensitivity analysis to be derived that can be applied to the trade-off between resource usage and schedulability (within the context of MCS).

### 3.4 Shared Resources

With mixed criticality systems it is not clear to what extent data should flow between criticality levels. There are strong objections to data flowing from low to high criticality applications unless the high criticality component is able to deal with potentially unreliable data [409] – this happens with some security protocols [83]. Even with data flowing in the other direction there remains the scheduling problem of not allowing a high criticality task to be delayed by a low criticality task that has either locked a shared resource for longer than expected or is executing at a raised priority ceiling level for too long.

Sharing resources within a criticality level is however a necessary part of any usable tasking model. In single criticality systems a number of priority ceiling protocols have been developed [39, 413]. These are beginning to be assessed in terms of their effectiveness for mixed criticality systems. Burns [94] extends the analysis for fixed priority systems by adding criticality specific blocking terms into the response-time analysis. He notes that the original form of the priority ceiling protocol (OPCP) [413] has some useful properties when applied to MCS. Resources can be easily partitioned between criticality levels and starvation of *LO*-criticality tasks while holding a lock on a resource can be prevented. With AMC-OPCP, a task can only suffer direct blocking if a resource is locked by a lower priority task of the same criticality.

Rather than use a software protocol, Engel [171] employs Hardware Transactional Memory to roll back any shared object to a previous state if a *LO*-criticality task overruns its budget while accessing the object.

For EDF-based scheduling Zhao et al. [475, 478] attempt to integrate the Stack Resource Protocol (SRP) [39] and Preemption Threshold Scheduling [460] with approaches to EDF scheduling that involve tasks having more than one deadline. This is not straightforward as these schemes assume that relative deadlines are fixed.

Alternative approaches are proposed by Lakshmanan et al. [292] by extending their single processor zero slack scheduling approach [356] to accommodate task synchronisation across criticality levels for fixed priority systems. They define two protocols: PCIP (Priority and Criticality Inheritance Protocol) and PCCP (Priority and Criticality Ceiling Protocol). Both of these contain the notion of criticality inheritance. This notion is also used by Zhao et al. [477] in their HLC-PCP (Highest-Locker Criticality Priority Ceiling Protocol) which they apply to the

AMC scheduling scheme (see Section 3.2.1). For a dual criticality system they define three modes of execution, the usual two plus an intermediate mode which covers the time during which *LO*-criticality tasks are allowed to continue to execute if they are holding a lock on a resource that is shared with a *HII*-criticality task.

A more systematic scheme is proposed by Brandenburg [90]. Here all shared resources are placed in *resource servers* and all access to these servers is via a MC-IPC protocol. As a result only these servers and the support for the MC-IPC protocol have to be developed to the highest criticality level. Resource users can be of any criticality level, including non-critical. Data sharing within the context of the *MC*<sup>2</sup> architecture (see Section 4) is addressed by Chisholm et al. [123].

### 3.5 Static and Synchronous Scheduling

The move between criticality levels can be captured in a static schedule by switching between previously computed schedules; one per criticality level. This is explored by Baruah and Fohler [68]. Socci et al. [419, 421, 422] show how these Time-Triggered (TT) tables can be produced via first simulating the behaviour one would obtain from the equivalent fixed priority task execution. Their approach is improved upon by Behera and Bhaduri [79] (their algorithm has lower computational complexity). Construction of the tables via tree search is addressed by Theis et al [437], and via the use of linear programming (LP) by Jan et al. [264]. For legacy systems Theis and Fohler [436] show how an existing single table may be used to support MCS.

A particularly simple table driven approach is to use a cyclic executive, this is investigated by Burns et al. [60, 99, 106, 151, 181, 184, 185] for multiprocessor systems in which the change from minor cycle to minor cycle is synchronised as is the change from executing code of one criticality to that of another. Both global and partitioned approaches are investigated, as are systems that use less processors for the *HII*-criticality work than they do for the *LO*-criticality work [185]. Both LP and ILP based formulations are used to construct the cyclic executives.

This use of tables is extended to synchronous reactive programs by Baruah [45, 46]. Here a DAG (Directed Acyclic Graph), of basic blocks that execute according to the synchrony assumption, is produced that implements a dual-criticality program. The synchronous approach is also considered by Yip et al. [466] and by Cohen et al. [132]. The latter proving an application of mixed criticality from the railway industry, and an example of why data needs to flow between criticality levels.

### 3.6 Varying Speed Processors

Most analysis for MCS assumes a constant speed processor, but there are situations in which the speed of the processor is not known precisely (for example with asynchronous circuitry). Baruah and Guo [70] consider power issues that could lead to a processor having variable speed. As a processor slows down the execution time of the tasks increase. They simplify the model by assuming two basic speeds, normal and degraded. At the normal speed a scheduling table is used; at the degraded speed only *HI*-criticality jobs are executed and they use EDF. The authors have extended this work [71, 219, 221] to include a more expressive model, issues of processor self-monitoring (or not), and a probabilistic approach to performance variation. They have also considered system which have both uncertainty in execution times and processor speed [222].

Voltage scheduling, and thereby variable speed computation, is used by Huang et al. [250, 251] to respond to a temporal overload – if a  $C(HI)$  value is exceeded and could lead to a *LO*-criticality task missing its deadline then energy is utilised to enable the processor to reduce computation times. Overall, their approach aims to reduce the system’s expected energy consumption. Taherin et al. [431] present some alternative DVFS schemes and compare their results with those of Huang. DVFS management is also addressed by Haririan and Garcia-Ortiz [230] in their provision of a simulation framework for power management.

## 4 Multiprocessor Analysis

The first paper to discuss mixed criticality within the context of multiprocessor or multi-core platforms was by Anderson et al. [20] in 2009 and then extended in 2010 [342]. Five levels of criticality were identified; going from level-A (the highest) to level-E (the lowest). They envisaged an implementation scheme, which they call  $MC^2$ , that used a cyclic executive (static schedule) for level-A, partitioned preemptive EDF for level-B, global preemptive EDF for levels C and D and finally global best-effort for level-E. They considered only harmonic workloads but allowed slack to move between containers (servers). Each processor had a container for each criticality level, and a two-level hierarchical scheduler (see Section 5.4). Later work from this group [124, 235] evaluates the OS-induced overheads associated with multiprocessor platforms. They also experimented with isolation techniques for LLC (last level cache) and DRAM. And have demonstrated, using  $MC^2$ , the benefits of having different isolation techniques for each criticality level [279, 280]. The support for mode changes within  $MC^2$  is considered by Chrisholm [122]. This  $MC^2$  framework is also used by Bommert [89] to support segmented mixed criticality parallel tasks.

In the remainder of this section we first look at task allocation (with global or partitioned scheduling), then consider analysis and finally communications and other systems resources. We note that there has also been work on implementing mixed-criticality synchronous systems on multiprocessor platforms [47].

#### 4.1 Task Allocation

The issue of allocation was addressed by Lakshmanan et al. [293] by extending their single processor slack scheduling approach [356] to partitioned multiprocessor systems employing a Compress-on-Overload packing scheme. Allocation in a distributed architecture was addressed by Tamas-Selicean and Pop [432–435] in the context of static schedules (cyclic executives) and temporal partitioning. They observed that scheduling can sometimes be improved by increasing the criticality of some tasks so that single-criticality partitions become better balanced. This increase comes at a cost and so they employ search/optimisation routines (Simulated Annealing [191, 433] and Tabu [256, 257, 432, 434]) to obtain schedulability with minimum resource usage. Search routines, this time GAs (Genetic Algorithms), are also used by Zhang et al. [473] to undertake task placement in security-sensitive MCS. Their objective is to minimise energy consumption “while satisfying strict security and timing constraints”. A toolset to aid partitioning is provided by Alonso et al. [17].

A more straightforward investigation of task allocation was undertaken by Kelly et al [277]. They considered partitioned homogeneous multiprocessors and compared first-fit and best-fit approaches with pre-ordering of the tasks based on either decreasing utilization or decreasing criticality. They used the original analysis of Vestal to test for schedulability on each processor, and concluded that in general first-fit decreasing criticality was best. For heterogeneous multiprocessors Awan et al. [35] propose a mapping scheme that is energy efficient.

A comprehensive evaluation of many possible schemes is reported by Rodriguez et al. [394]. They consider EDF scheduling and used the analysis framework of EDF-VD (see Section 3.3). One of their conclusions is the effectiveness of a combined criticality-aware scheme in which *HI*-criticality tasks are allocated Worst-Fit and *LO*-criticality tasks are allocated using First-Fit; both with Decreasing Density. The same result is reported by Gu et al. [209]. They additionally note that if there are some very ‘heavy’ *LO*-criticality tasks (i.e. high utilisation or density) then space must be reserved for them before the *HI*-criticality tasks are allocated. Partitioning with EDF-VD is also addressed by the work of Han et al. [227]. Ramanathan and Easwaran [385] demonstrate that an effective partitioning scheme can be derived from evenly distributing the differences between *HI*-criticality and *LO*-criticality utilisation. Their results being applicable to EDF-VD

and fixed-priority AMC.

A global allocation scheme for MCS is proposed by Gratia et al. [201, 203]. They adapt the RUN scheduler [392], which uses a hierarchy of servers, to accommodate *HI* and *LO* criticality tasks. The latest version of their scheduler (GMC-RUN) [202] has been extended to deal with more criticality levels.

Between fully partitioned and fully global scheduling is the class of schemes termed *semi-partitioned*. This is being addressed by Bletsas et al. [32, 86, 87] and Al-Bayati et al. [8]. The latter work uses two allocations for their two criticality modes. *HI*-criticality tasks do not migrate. During a mode change, carry-over *LO*-criticality jobs are dropped and new *LO*-criticality jobs executing on a different processor are given extended deadlines/periods (i.e. they utilise the elastic task model). A different approach is taken by Xu and Burns [464]; here a mode change on one processor results in *LO*-criticality jobs migrating to a different processor that has not suffered a criticality mode change. No deadlines are missed. If all processors suffer such a mode change then at least the timing needs of all *HI*-criticality tasks are protected.

A different approach, that aims to maximise the benefits of partitioning and global scheduling, is that based on clusters. A multi-core platform is statically partitioned into a number of clusters, within a cluster tasks execute ‘globally’. Ali and Kim [13] investigate a scheme in which small clusters are used when the system is in the *LO*-criticality mode, but larger clusters are employed when the system moves to the *HI*-criticality mode. Nagalakshmi and Gomathi [347] also use clusters but within a EDF-scheduled fully partitioned approach. Within each cluster is a single *HI*-criticality task and one or more *LO*-criticality tasks. As a result the impact of a *HI*-criticality task executing beyond its  $C(LO)$  value is limited to just those *LO*-criticality tasks within the same cluster.

With dual-criticality fault tolerant systems, a scheme in which high criticality tasks are replicated (duplicated) while low criticality tasks are not is investigated by Axer et al. [36] for independent periodic tasks running on a MPSoC (multi-processor system-on-chip). They provide reliability analysis that is used to inform task allocation.

A more theoretical approach (i.e. it is not directly implementable) is proposed by Lee et al. [298] with their MC-Fluid model. A fluid task model [65, 241] executes each task at a rate proportional to its utilisation. If one ignores the cost of slicing up tasks in this way then the scheme delivers an optimal means of scheduling multiprocessor platforms. To produce a mixed criticality version of the fluid task model the fact that tasks do not have a single utilisation needs to be addressed. Lee et al. [298, 299] do this and they also produce an implementable version of the model that performs well in simulation studies (when compared with other approaches). Baruah et al. [41, 66] derived a simplified fluid algorithm which they

call MCF. Two further algorithms, MC-Sort and MC-slope, are proposed by Ramanathan and Easwaran [384, 387].

All the above work is focussed on standard single threaded tasks. In addition there has been some studies on parallel tasks and MCS – see Liu et al. [309, 310, 322], Pathan [373] and Agrawal and Baruah [4].

## 4.2 Schedulability Analysis

For globally scheduled systems Li and Baruah [308] take a ‘standard’ multiprocessor scheme, fpEDF [42] and combine it with their EDF-VD approach (see Section 3.3). Evaluations indicate that this is an effective combination. Extensions of this work [64] compare the use of partitioning or global scheduling for MCS. Their interim conclusion is that partitioning is by far the most effective approach to adopt.

Notwithstanding this result, Pathan derives [370] analysis for globally scheduled fixed priority systems. They adopt the single processor approach [61] (see Section 3.2) and integrate this with a form of analysis for multiprocessor scheduling that is amenable to optimal priority ordering, via Audsley’s algorithm [26]. They demonstrate the effectiveness of their approach (by comparing success ratios). Jung and Lee subsequently improved on this analysis [270].

A different and novel approach to multi-core scheduling of MCS is provided by Kritikakou et al. [285, 288]. They identify that a *HI*-criticality task will suffer interference from a *LO*-criticality task running on a different core due to the hardware platform’s use of shared buses and memory controllers etc.. They monitor the execution time of the *HI*-criticality task and can identify when no further interference can be tolerated. At this point they abort the *LO*-criticality task even though it is not directly interfering. An implementation on a multi-core platform demonstrated effective performance of their scheme [288]. They subsequently improved on this static approach by utilising a dynamic version that reduces the time spent in the controller [286].

Extensions to deal with precedence constraints were given by Socci et al. [420] but only for jobs (not tasks). A full pipeline scheme is considered by de Niz et al. [150].

## 4.3 Communication and other Resources

With a more complete platform such as a multiprocessor or System on Chip (SoC), perhaps with a NoC (Network-on-Chip), more resources have to be shared between criticality levels. The first design issue is therefore one of partitioning (as addressed above), how to ensure the behaviour of low criticality components does not adversely impact on the behaviour of higher criticality components. Pellizzoni

et al. [376] in 2009 was the first to consider the deployment of mixed criticality systems (MCS) on multi-core and many-core platforms. They defined an Architectural Analysis and Design Language (AADL), a form of ADL (Architectural Description Language), for mixed criticality applications that facilitates system monitoring and budget enforcement of all computation and communication. Later Obermaisser et al. [363, 364] introduce a system model with gateways and end-to-end channels over hierarchical, heterogeneous and mixed criticality networks.

For a bus-based architecture it is necessary to control access to the bus so that applications on one core do not impact unreasonably on applications on other cores (whether of different or indeed the same criticality level). Pellizzoni et al. [377] show that a task can suffer a 300% increase in its worst-case execution time due to memory access interference even when it only spends 10% of its time on fetching from external memory on an 8-core system. To counter this, Yun et al. [467] propose a memory throttling scheme for MCS. Kotaba et al [284] also propose a monitoring and control protocol to prevent processes flooding any shared communication media be it a bus or network. Kritikakou et al. [287] consider a scenario in which there are a few critical tasks that can suffer indirect interference from many lower critical tasks. They attempt to allow as much parallelism as possible commensurate with the critical tasks retaining their temporal validity. Hassan and Patel [231] claim an improved bus arbitrator, called Carb, that is more criticality aware. Bounding the interference that a safety-critical task can suffer from lower criticality tasks using the same shared communication resources on a multi-core platform is also addressed by Nowotsch et al. [361]. Freitag et al. [187] utilise a fully isolated model (each core only has tasks of only one criticality), those core that have *LO*-criticality tasks are slowed down (or even stopped) if their measured interference on the *HI*-criticality cores is above statically derived bounds.

Within the time-triggered model of distributed computation and communication a mixed criticality system is often viewed as one that has both time-triggered and event-triggered activities, also referred to as synchronous and asynchronous [382, 425]. The time-triggered traffic is deemed to have the highest criticality, the event-triggered traffic can be either just best-effort or can have some level of assurance if its impact on the system is bounded; what Steiner [425] calls *rate-constrained*. Protocols that support this distinction can be supported on networks such as TTEthernet. Another TDMA-based approach, though this time built into the Real-Time Ethernet protocol, is proposed by Carvajal and Fischmeister in their open-source framework, *Atacama* [114]. Cilku et al. [128] describe a TDMA-based bus arbitration scheme. Novalk et al. [359] propose a scheduling algorithm for time-triggered traffic that minimises jitter while allowing *HI*-criticality messages to be re-transmitted (following failure) at the expense of *LO*-criticality messages (which are abandoned). They also [360] consider how to produce an effective

static schedule when there are unforeseen re-transmissions (for two and three levels of criticality). Analysis of AFDX (Avionics Full Duplex Switched Ethernet) within the context of MCS is provided by Finzi et al. [180].

A reconfigurable SDRAM controller is proposed by Goossen et al. [199] to schedule concurrent memory requests to the same physical memory. They also use a TDMA approach to share the controller's bandwidth. A key aspect of this controller is that it can adapt to changes in the run-time characteristics of the application(s). For example, a criticality mode change which should result in more bandwidth being assigned to the higher criticality tasks can be accommodated by what the authors call a *use-case switch*. Criticality aware DRAMs are also addressed by Jalle et al [263] in the context of a Space case study in which there are two criticality levels: 'control' and 'payload'.

Virtual DRAMs are adapted by Ecco et al. [163] to isolate critical tasks (which are guaranteed) from non-critical tasks that, although not guaranteed, do perform adequately. Each virtual device represents a group of DRAM banks supports one critical task and any number of non-critical tasks. All critical tasks run on dedicated cores, and hence the only potential source of inter-criticality interference is from the interconnection fabric (bus). By use of virtual devices, the critical tasks benefit from interference-free memory access. DRAMs are also the focus of the work by Hassen et al [232], Awan et al [31] and Guo and Pellizzon [218].

Kim et al. [278] propose a priority-based DRAM controller for MCS that separates critical and non-critical memory accesses. They demonstrate improved performance for the non-critical traffic. Note this work is focussed on supporting critical and non-critical traffic on the same memory banks (rather than mixed-criticality). A similar approach and result is provided by Goossens et al. [198] with their open-page policy.

Giannopoulou et al. [191, 192] use a different time-triggered approach. They partition access to the multiprocessor bus so that at any time,  $t$ , only memory accesses from tasks of the same criticality can occur. This may introduce some inefficiencies, but it reduces the temporal modelling of a mixed criticality shared bus to that of a single criticality shared bus. The latter problem is not, however, straightforward (but is beyond the scope of this review). In later work they generalise their approach by introducing the notion of isolation scheduling [247]. They also demonstrate the implementation of their approach on a Layray MPPA-256 many-core platform [446].

The problems involved in using a shared bus has lead Giannopoulou et al. to also include a Network-on-Chip (NoC) in their later work [194]. Burns et al. [60, 99, 106, 184] apply a 'one criticality at a time' approach to MCS scheduled by the use of a Cyclic Executive; they considered both partitioned and global allocation of jobs to frames.



Tobuschat et al. [444, 445] have developed a NoC explicitly to support MCS. Their IDAMC protocol uses a *back suction* technique [153] to maximise the bandwidth given to low (or non) critical messages while ensuring that high-criticality messages arrive by their deadlines. The more familiar wormhole routing [355] scheme for a NoC has been expanded by Burns, Harbin and Indrusiak [107, 258] to provide support for mixed criticality traffic. Response-time analysis, already available for such protocols [414]<sup>21</sup>, is augmented to allow the size and frequency of traffic to be criticality aware. Wormhole routing is also used by Hollstein et al [240] to provide complete separation of mixed-criticality code; they also support run-time adaptability following any fault identified by a Built-In Self Test. Another wormhole router (DAS - double arbiter and switching) is described by Dridi et al. [158, 159]; they use wormhole for the *LO*-criticality traffic, but store-and-forward for the *HI*-criticality messages.

On-chip networks require reliable/trusted interfaces to prevent babbling behaviour [92]; Ahmadian and Obermaisser [6] describe how to provide this via a time-triggered extension layer for a mixed-criticality NoC. Dynamic and adaptive control of a mixed-criticality NoC is considered by Kostrzewa et al. [282, 283]. Control over I/O contention via an Ethernet-based criticality-aware NoC is advocated by Abdallah et al. [1]. A focus on NoC security, in which *HI*-criticality messages need more protection than *LO*-criticality is taken by Papastefanakis et al. [368].

An alternative to using a NoC for all traffic (task to task and task to off chip memory) is proposed by Audsley [27, 196]. They advocate the use of a separate memory hierarchy to link each core to off chip memory. A criticality aware protocol is used to pass requests and data through a number of efficient multiplexers. If the volume of requests and data is criticality dependent then analysis similar to that used for processor scheduling can be used on this memory traffic. The separation of execution-time from memory-access time is explored by Li et al. [312, 313]. They demonstrate that this distinction improves schedulability.

Controller Area Network (CAN) [195] is a widely used network for real-time applications, particularly in the automotive domain. It has been the subject of considerable attention with Response-Time Analysis derived [146] for what is, in essence, a fixed priority non-preemptive protocol. The use of CAN in mixed criticality applications has been addressed by Burns and Davis [101]. In this work it is the period of the traffic flows and the fault model that changes between criticality levels. A MixedCAN protocol was developed that makes use of a Trusted Network Component that polices the traffic that nodes are allowed to send over the

---

<sup>21</sup>This analysis has been shown to be optimistic in some circumstances, see discussion by Xiong et al. [463].

network. Evaluations are used to show the advantages of using MixedCAN rather than a criticality agnostic approach. However the paper, in keeping with many other publications, only considered dual-criticality systems.

Herber et al. [234] also addressed the CAN protocol. They replaced the physical network controller with a set of virtual controllers that facilitate spacial separation. A weighted round robin scheduler is then used to give temporal isolation. Their motivation is to support virtualisation in an automotive platform. They do not however use criticality specific parameters for the different applications hosted on the same device. Nager et al. [348] show how CAN can be used in a more flexible way to support MCS.

Other protocols that have been considered in terms of their support for mixed criticality systems include FlexRay [200] and switched Ethernet [134, 135]. In the latter work, a change in criticality mode is broadcast to the entire system by adding a new field to the IEEE 1588 PTP (Precision Time Protocol). George et al. [188] assume the speed of the wireless communication media varies over time. Each node monitors this speed and only send messages that are compatible with its current estimate. Offline analysis partitions the messages according to criticality and required speed.

A further communication protocol is addressed by Addisu et al. [3]. They consider JPEG2000 Video streaming over a wireless sensor network. With such a network the available bandwidth varies in an unpredictable way. They propose a bandwidth allocation scheme that is criticality aware. A wireless protocol (WirelessHART) is also used by Jin et al. [266, 267] to support delay analysis with fixed priority scheduling for sensor networks. A hybrid scheduling approach for sensor nodes with mixed-criticality tasks is investigated by Micea et al. [339].

A means of modelling an entire end-to-end vehicular embedded system including various forms of networking is provided by Mubeen et al. [344].

A novel scheduling approach (triangle scheduling) for mixed-criticality messages is proposed by Dürr et al. [160]. The adaptation of synchronisation protocols such as MSRP (Multiprocessor Stack Resource Policy) to multiprocessor platforms and mixed-criticality software scheduled using partitioned EDF is addressed by Han et al. [226].

## **5 Links to other Research Topics**

### **5.1 Hard and Soft Tasks**

Although the label ‘Mixed Criticality Systems’ is relatively new, many older results and approaches can be reused and reinterpreted under this umbrella term. In particular dual-criticality systems in which there are hard and soft tasks combined

has been studied since at least 1987 [302]. Hard tasks must be guaranteed. Soft tasks are then given the best possible service. But soft tasks are usually unbounded in some sense (either in terms of their execution time or their arrival frequency) and hence they must be constrained to execute only from within *servers* (execution-time servers). Servers have bounded impact on the hard tasks. Since 1987 a number of servers have been proposed. The major ones for fixed priority systems being the Periodic Server, the Deferrable Server, the Priority Exchange Server (all described by Lehoczky et al. [302]), and the Sporadic Server [423]. The ability to run soft tasks in the slack provided by the hard tasks is also supported by the Slack Stealing schemes [139, 148, 301, 389] which have similar properties to servers. These all have equivalent protocols for dynamic priority (EDF) systems; and some EDF specific ones exist such as the Constant Bandwidth server [316].

Since their initial specification, analysis has improved and means of allocating and sharing capacity between servers have been investigated (see for example [80, 81, 468]). However, these results on servers (and examples of how they can be implemented in Operating Systems and programming languages) are generally known (see standard textbooks [109, 112, 323]), hence they are not reviewed in more detail here. Note however, that standard servers only deal with the isolation/partitioning aspect of MCS. To support sharing (of resources) there must be some means of moving capacity from the under utilised servers of high criticality tasks to the under provisioned servers of lower criticality tasks. The Extended Priority Exchange server [423] as well as work on making use of gain time, show how this can be achieved.

Another way of maximising the slack available for soft tasks is the dual-priority scheme [110, 149]. Here there are three bands of priority. The soft tasks run in the middle band while the hard tasks start in the lower band but are promoted to the higher band at the latest possible time commensurate with meeting their deadlines. So hard tasks execute when they have to, or when there are no soft tasks, soft tasks run otherwise.

Run-time adaptability for MCS has been addressed by Hu et al. [242, 243]. They present an approach to adaptively shape at runtime the inflow workload of *LO*-criticality tasks based on the actual demand of *HI*-criticality tasks. This improves the QoS of *LO*-criticality tasks; but it not clear what level of guarantee is provided for these tasks. An alternative scheme, with the same aim, is given by Hikmet [238]. A QoS focus is also taken by Vaidhun et al. [449] with respect to pervasive systems.

## 5.2 Fault Tolerant Systems (FTS)

Fault tolerant systems typically have means of identifying a fault and then recovering before there is a system failure. Various recovery techniques have been proposed including exception handling, recovery blocks, check-points, task re-execution and task replication. If, following a fault, extra work has to be undertaken then it follows that some existing work will need to be abandoned, or at least postponed. And this work must be less important than the tasks that are being re-executed. It follows that many fault tolerant systems are, in effect, mixed criticality.

To identify a fault, timeouts are often used. A job not completing before a deadline is evidence of some internal problem. Earlier warning can come from noting that a job is executing for more than its assumed worst-case execution time. Execution-time monitoring is therefore common in safety critical systems that are required to have at least some level of fault tolerance. Again this points to common techniques being required in FTS and MCS.

As noted earlier, in the discussion on CAN (Section 4.3), a fault model can be criticality dependent [101] – a task may, for example, be required to survive one fault if it is mission critical, but two faults if it is safety critical. The difference between the assumed computation times at different criticality levels may be a result of the inclusion or not of recovery techniques in the assumed worst-case execution time of tasks.

Although there is this clear link between FTS and MCS there has not yet been much work published that directly addresses fault-tolerant mixed criticality systems. Exceptions being work by Huang et al. [253, 254], a paper by Pathan [371] that focuses on service adaptation and the scheduling of fault-tolerant MCS, and a four-mode lockstep model developed by Al-Bayati et al. [7] (the requirement for lockstep execution is relaxed in later work [113]). Work by Thekkilakattil uses Zonal Hazard Analysis and Fault Hazard Analysis [439] and Error-Burst models [438] to deliver both flexibility and real-time guarantees for the most critical tasks. Thekkilakattil et al. [440] also considers the link between MCS and the tolerance of permanent faults. Lin et al. [314] attempt to integrate mixed criticality with the use of primary and backup executions in both of the two criticality modes they consider. Islam et al. [259], in a paper that preceded that of Vestal, looked at combining different levels of replication for different levels of criticality. Co-scheduling loosely-coupled replicas for MCS is shown to be effected by Rambo and Ernst [388]. Ahmadian et al. [5] invoke reconfiguration to provide for fault tolerance in the context of a time-triggered NoC based MCS. Alahmad and Gopalakrishnan [10] use what they term ‘isochronous’ execution to synchronise the execution of redundant versions of critical software over ‘regular’ and ‘safe’ processors.

As highlighted already in this review, many models and protocols for mixed criticality behaviour allow the system to move through a sequence of criticality modes. With a two mode system (*HI* and *LO*) the system starts in the *LO* mode in which all deadlines of all tasks are guaranteed, but can then transition to the *HI* mode in which only the *HI*-criticality tasks are guaranteed (and the *LO*-criticality tasks may actually be abandoned). It may, or may not, later return to the *LO* mode when it is safe to do so. Burns [95] attempts to compare these criticality mode changes with the more familiar system mode change. He concludes that the *LO* mode behaviour should be considered to be the ‘normal’ expected behaviour. A move away from this mode is best classified as a fault; with all other modes being considered forms of *graceful degradation*. Such a mode change is *planned but may never occur*. A move back to the fully functional *LO* mode is closest in nature to an *operational* (sometime known as exceptional) mode change.

### 5.3 Security

Many MCS papers have, either explicitly or implicitly, focused on issues of safety and reliability. Criticality can however also refer to security. Within this domain it is usual to have different security levels. And hence much of the extensive literature on security is relevant, but is out of scope of this review. Some work is nevertheless applicable to safety and security; for example the definition of a *separation kernel* for a system-on-chip built using a time-triggered architecture [461]. An overview of the security (and other) issues associated with MPSoCs is provided by Hassan [233].

Another paper directly linking security and MCS is from Baek and Lee [38]. They incorporate the cost of flushing between security levels in a non-preemptive version of AMC (for fixed priority scheduling).

### 5.4 Hierarchical Scheduling

One means of implementing a MCS where strong partitioning is needed between applications is to use a hierarchical (typically two-level) scheduler. A trusted base scheduler assigns budgets to each application. Within each application a secondary scheduler manages the threads of the application. There are a number of relevant results for such resource containment schemes (e.g. [117, 143, 144, 317, 396, 470]). Both single processor and multiprocessor platforms can support hierarchical scheduling.

Unfortunately when hierarchical scheduling is applied to MCS there is a loss of performance [291]. A simple interface providing a single budget and replenishment period (which is often associated with virtualisation or the use of a hypervi-

sor [16]) is too inflexible to cater for a system that needs to switch between criticality levels. To provide a more efficient scheme, Lackorznschi et al. [291] propose ‘flattening’ the hierarchy by exposing some of the interval structure of the scheduled applications. They develop the notion of a scheduling context which they apply to MCS [455]. In effect they assign more than one budget to each ‘guest’ OS. As a result, applications that would otherwise not be schedulable are shown to utilise criticality to meet all deadlines. An alternative, but still flexible approach, is provided by Groesbrink et al. [206, 207]. They allow budgets to move between virtual machines executing on a hypervisor that is itself executing on a multi-core platform. The hypervisor controls access to the processor, the memory and shared I/O devices. Yet another scheme is described by Marinescu et al [328]; they are more concerned with partitioning as opposed to resource usage, but they do address distributed heterogeneous architectures. Hypervisors are also used by Cilku and Puschner [129], to give temporal and spacial separation on a multiprocessor platform, and Perez et al. [379] use a hierarchical scheduler to statically partition a wind power mixed criticality embedded system requiring certification under the IEC-61508 standard. A hypervisor for a mixed criticality on-board satellite software system is discussed by Salazar et al. [15, 399] and one for general control system is addressed by Crespo et al [133]. The issue of minimising the overheads of a hypervisor is addressed by Blin et al. [88].

## 5.5 Cyber Physical Systems and Internet of Things

In parallel with the development of a distinct branch of research covering MCS has been the identification of Cyber Physical Systems (CPS) as a useful focus for system development. Not surprisingly it has been noted that many CPS are also mixed criticality. For example Schneider et al. [406] note that many CPS are a combination of deadline-critical and QoS-critical tasks. They propose a layered scheme in which QoS is maximised while hard deadline tasks are guaranteed. Izosimov and Levholt [262] use a safety-critical CPS to explore how metrics can be used to map potential hazards and risk from top level design down to mixed criticality components on a multi-core architecture. Issues of composability within an open CPS are introduced in the short paper by Lee et al. [296].

Maurer and Kirner [332] consider the specification of cross-criticality interfaces (CCI) in CPSs that define the level of communication allowed between ‘open’ subsystems/components. Lee et al. [295] also look at interfaces and composition for mixed criticality CPSs.

The link between the Internet of Things, IoT, and MCS is made by Kamienski et al. [276] in the context of development methods for energy management in public buildings. Smart buildings are also the focus of the work of Dimopoulos et al. [155]




Figure 2: pWCET distribution as an Exceedance function

on a context-aware management architecture.

## 5.6 Probabilistic real-time systems

In mixed criticality systems, the worst-case execution time of a task is expressed as a function of the criticality level (e.g.  $C(LO)$  and  $C(HI)$ ) with larger values for the WCET obtained for higher criticality levels. Research into probabilistic hard real-time systems can be viewed as extending this model to a continuum (or at least a large number of discrete values). Instead of a number of single values for the WCET with different levels of confidence, the worst-case execution time is expressed as a probability distribution, referred to as a pWCET [82].

The exceedance function (or  $1 - \text{CDF}$ <sup>22</sup>) for the pWCET gives the probability that the task will exceed the specified execution time budget on any given run. Conversely, the exceedance function may be used to determine the execution time budget required such that the probability of overrunning that budget does not exceed a specified probability. This is illustrated in Figure 2. Here, an execution time budget of 55 has a probability of being exceeded of  $10^{-5}$ , whereas the execution time budget required to ensure that the probability of exceedance is at most  $10^{-9}$  is 70. We note that exceedance probabilities and failure rates (e.g.  $10^{-9}$  fail-

<sup>22</sup>Cumulative Distribution Function.

ure per hour) are not the same, but that such probabilities can be transformed into failure rates by accounting for the number of jobs in a given time period, or via probabilistic schedulability analysis techniques.

Probabilistic analysis provides an alternative treatment for mixed criticality systems, where high criticality tasks are specified as having an extremely low acceptable failure rate (e.g.  $10^{-9}$  per hour), whereas a higher failure rate (e.g.  $10^{-6}$  or  $10^{-7}$  per hour) is permitted for lower criticality tasks. Probabilistic worst-case execution times [18, 115] and the probabilistic worst-case response times [152, 324] derived from them provide a match to requirements specified in this way. These techniques can potentially be used to show that pathological cases with very high execution times / high response times have a provably vanishingly low probability of occurring, thus avoiding the need to over-provision compute resources to handle these cases.

Just as MCS has expanded from a focus on worst-case execution times to one that includes arrival rates (for sporadic work), probabilistic analysis has been developed [12] for the case where the arrival rate of tasks is described by a probability distribution. This work could form a further link between MCS and probabilistic analysis. Indeed Masrur [329] uses random jitter on the arrival time of *LO*-criticality tasks to improve schedulability.

Guo et al. [223] demonstrate the usefulness of a probabilistic framework in their analysis of an EDF scheduled system in which there is a permitted (but low) probability of timing faults. The chances of a *HI*-criticality task executing for more than its *LO*-criticality value is also expressed as a probability. Their current work assumes that task execution times are independent; this is an unrealistic assumption, but one that could be weakened in future work. Santinelli and George [401] also explore the probability space of worst-case execution times for MCS. Probabilistic analysis for the SMC and AMC schemes is derived by Maxim et al. [333, 334]. And a constrained Markov decision process is used by Alahmad and Gopalakrishnan [9, 11] to model job releases in MCS. Probabilistic analysis is also used to investigate the safety of each criticality level [157]. Abdeddaim and Maxim [2] derived probabilistic response time analysis for mixed criticality tasks under fixed priority preemptive scheduling, computing the probability of deadline misses for each task in each criticality mode. Their work does not assume any monitoring, hence they assume that lower criticality tasks continue to execute in higher criticality modes. Kttler et al. [290] use symbolic execution to derive probabilistic estimates of the *LO*-criticality tasks' completing before their deadlines.

One further use of stochastic models is to estimate the amount of time that *LO*-criticality tasks are unavailable (not executed). Medina et al. [335, 337] provide such availability analysis for MCS hosted on multi-core platforms scheduled using data-flow graphs.



## 6 More Realistic MCS Models

The abstract behavioural model described in Section 2 has been very useful in allowing key properties of mixed criticality systems to be derived, but it is open to criticism from systems engineers that it does not match their expectations. In particular:

- In the *HI*-criticality mode, *LO*-criticality tasks should not be abandoned. Some level of service should be maintained if at all possible, as *LO*-criticality tasks are still critical.
- For systems which operate for long periods of time it should be possible for the system to return to the *LO*-criticality mode when the conditions are appropriate. In this mode all functionality should be provided.

It can be argued that these criticisms are, at least partly, misplaced as any high integrity system should remain in the *LO*-criticality mode for its entire execution: the transition to *HI*-criticality mode is only a theoretical possibility that the scheduling analysis can exploit [61]. Nevertheless, in less critical applications (such as those envisaged in the automotive industry) actual criticality mode changes may be experienced during operation and the above criticisms should be addressed. Of course for some applications it is acceptable to provide only limited timing guarantees during these rare events, and hence no online controls are required [457].

Recent reflections (at two 2017 events – the WMC workshop and Dagstuhl seminar) has identified two distinct roles for what is now called the ‘Vestal Model’: the verification of a MCS, and its run-time survivability. For the former it may be only necessary to argue that the *HI*-criticality work is guaranteed (in all circumstances) and hence no guarantees need be developed for the *LO*-criticality work. But survivability concerns all criticality levels. If faults are occurring (in particular timing overruns) then criticality-aware graceful degradation should result. Note material linking MCS and fault tolerance was reviewed in Section 5.2.

With this dual role for the Vestal Model it is not true to say that it advocates dropping all *LO*-criticality tasks as soon as there is any, even minor, budget overrun. Rather it is saying that for static verification the interference from lower criticality tasks must be demonstrably bounded. Within this bound a variety of techniques are available to increase the survivability of the *LO*-criticality work.

To distinguish between two forms of survivability, Burns et al. [100] define Robustness to imply those techniques used to deliver full functionality within an explicit bound on the duration and severity of the temporal faults, and Resilience to

encompass the wide range of techniques available to give different levels of graceful degradation when the system’s faulty behaviour is beyond the lower bound. These techniques include:

1. Letting any *LO*-criticality job that has started, run to completion (this is in effect what is assumed by many forms of analysis [61]).
2. Reducing the priorities of the *LO*-criticality tasks [58], or similar with EDF scheduling [248, 249].
3. Increasing the periods and deadlines of *LO*-criticality jobs [265, 386, 426, 428–430], called *task stretching*, the *elastic task model* or *multi-rate*.
4. Imposing only a weakly-hard constraint on the *LO*-criticality jobs [189, 336].
5. Decreasing the computation times of some or all of the *LO*-criticality tasks [98]<sup>23</sup> [41], perhaps by utilising an imprecise mixed-criticality (IMC) model [320, 321, 372] or budget control [212].
6. Moving some *LO*-criticality tasks to a different processor that has not experienced a criticality mode change [464].
7. Abandoning *LO*-criticality work in a disciplined sequence [183, 215, 252, 297, 393].

The fifth action leads to a modification to the system model; whereas for *HI*-criticality tasks we have  $C(HI) \geq C(LO)$ , for *LO*-criticality tasks we now have  $C(HI) \leq C(LO)$ . For some tasks  $C(HI) = 0$ , that is they are abandoned. For others a lower level of service can be guaranteed. For some they may be able to continue with their full computation time budgets.

The final approach is addressed by Fleming and Burns [183]; they introduce a further notion into the standard model; tasks are allocated to *applications* and each application is assigned (by the system designers) an *importance* level. *LO*-criticality tasks are abandoned in inverse order of importance. Huang et al. [252] also introduce an extension to the standard model by the use of an ICG (Interference Constraint Graph) to capture more specifically which tasks need to be dropped when particular higher criticality tasks exceed their allocated criticality-aware execution times. Controlled abandonment by the use of partitioning is advocated by Mahdiani and Masrur [327] in the context of the EDF-VD scheduling.

Obviously all seven schemes can be used together: complete or move all started jobs, allow some new jobs to have an extended deadline or reduced computation

<sup>23</sup>Note equation (6) in this paper has a typo, both  $R_i(LO)$  terms should be starred ( $R_i^*(LO)$ ).

time or a weakly-hard constraint, reduce the priorities of some others jobs and abandon those of lowest importance in particular partitions. Such an approach is advocated by Su et al. [427]; here *LO*-criticality tasks have two periods (short and long) and two priorities. At the criticality mode change these tasks switch to their longer periods and new priorities. Analysis is provided to show that all modes are schedulable.

A flexible scheme utilising hierarchical scheduling is proposed by Easwaran et al. [162, 214]. They differentiate between minor violations of *LO*-criticality execution time which can be dealt with within a component (an internal mode change) and more extensive violations that requires a system-wide external mode change. In doing so they introduce a new mixed-criticality resource interface model for component-based systems which supports isolation, virtualisation and compositionality.

Within the context of EDF-VD Chen et al. [118] allow the overrunning *HI*-criticality task to progress through multiple intermediate levels rather than assume a single jump to  $C(HI)$ . They term this a Multiple-Shot transition. It facilitates a more control impact on *LO*-criticality tasks. A similar idea is utilised by Chwa et al. [126].

Papadopoulos et al. [367] also address small overruns of *LO*-criticality budgets, they use a control-theoretic approach to automatically make minor and stable modifications to future budgets so as to return the system to normal behaviour without the need for mode changes or other significant reductions in system functionality. A more controlled response to small overruns is also considered by Massaro et al. [331]; they use a proactive mode change to anticipate an overrun before it actually occurs, thereby reducing its impact.

In keeping with other mode change situations ( [108, 170, 375, 390, 412, 442, 443]) a simple protocol for controlling the time of the change of mode back to *LO*-criticality is to wait until the system is idle (has no application tasks to run) and then the change can safely be made [442]. Santy et al. [405] extend this approach and produce a somewhat more efficient scheme that can be applied to globally scheduling multiprocessor systems, in which the system may never get to an idle tick. With a dual criticality system that has just transitioned into the *HI*-criticality mode (and hence no low-criticality jobs are executing); their protocol first waits until the highest priority *HI*-criticality job completes, then it waits until the next highest priority job is similarly inactive. This continues until the lowest priority job is inactive; it is then safe to reintroduce all *LO*-criticality tasks. Obviously if there is a further violation of the  $C(LO)$  bound then the protocol is abandoned and subsequently restarted. The authors call this a SCR (Safe Criticality Reduction); their paper also has a second protocol, but this is less intuitive and considerably more expensive at run-time.

A more aggressive scheme for returning a system back to its  $LO$ -criticality mode is proposed by Bate et al. [77, 78]. In this approach a bailout protocol is proposed.  $HI$ -criticality tasks take out a loan if they execute for more than their  $C(LO)$  estimate. Other tasks repay the loan by either not executing at all or by executing for less than expected. When the loan is repaid (and a further condition is met) the system returns to its normal mode. The authors demonstrate, using a scenario-based assessment, that the bailout protocol returned the system to the normal mode much quicker than the ‘wait for idle tick’ scheme.

As well as experiencing a criticality mode change a system can, of course, be structured to behave in a number of operational or behavioural modes. As indicated earlier, Burns [95] compares and contrasts these two forms of mode change. De Niz and Phan [357] note that the criticality of a task can depend on the behavioural mode of the system. They develop scheduling analysis for this dependency and consider the static allocation of such tasks to multiprocessor platforms.

Another aspect of the ‘standard model’ for MCS that can be argued to be unrealistic is the idea that a system with, say, five criticality levels would also have five different estimates of worst-case execution time for its most critical tasks. An augmented model has been proposed [96, 356] that restricts each task to having just two estimates of WCET. So, in the general case where there are  $V$  criticality levels,  $L_1$  to  $L_V$  (with  $L_1$  being the highest criticality), each task just has two  $C$  values. One represents its estimated execution time at its own criticality level ( $C_i(L_i)$ ) and the other an estimate at the base (i.e., lowest) criticality level ( $C_i(L_V)$ ). It follows that if a job is of the lowest criticality level (i.e.,  $L_i = L_V$ ) then it only has one WCET parameter. For all other jobs,  $C(L_i) \geq C(L_V)$ . The two parameters of this augmented model have been referred to as  $C(self)$  (or  $C(SF)$ ) and  $C(normal)$  ( $C(NL)$ ); the model seems to be sufficiently expressive to capture most of the key properties of mixed criticality systems. However, Baruah and Guo [72] has shown that: “*The Burns model is strictly less expressive than the Vestal model. Determining whether a given instance can be scheduled correctly remains NP-hard in the strong sense. Lower bounds on schedulability, as quantified using the speedup factor metric, are no better for the Burns model than for the Vestal model.*” This quote makes use of terms described in the next section.

An important property of any realistic form of analysis is that it is sustainable [57]. This, informally, implies that a system that is deemed schedulable by some appropriate test will remain schedulable when the application’s characteristics improve (for example, worst-case execution times are reduced, or period and deadlines extended). Guo et al [224] consider the sustainability of various forms of analysis for MCS. For example, they showed that AMC and EDF-VD are sustainable forms of analysis. This not only included reductions in  $C(HI)$  and  $C(LO)$  but also a reduction in the criticality level assigned to a task (e.g. a reduction

from *HI*-criticality to *LO*-criticality). They showed that other methods of implementing MCS such as criticality-monotonic are not sustainable. Sustainability of MCS is also considered by Kahil et al [271, 273] who introduce the notion of MC-Sustainable, and link sustainability to predictability.

## 7 More Formal Treatments

In this section we consider utilisation bounds, speedup factors and (formal) language issues.

### 7.1 Utilisation Bounds

For normal single criticality systems there are well known bounds on task set utilisation that will deliver a schedulable system with either fixed priority or EDF scheduling. Although the definition of utilisation is not straightforward when a task has more than one worst-case computation time, it is possible to give an effective definition and to derive least upper bounds (LUBs) for MCS. Santos-Jr et al. [403] derive a number of useful results for LUB. They construct a task set that is unschedulable (during a criticality mode change) with LUB arbitrarily close to 0. But where tasks have harmonic periods LUB can reach 1 (for a uniprocessor system). Between these two extremes they show that if higher criticality tasks do not have periods longer than lower criticality tasks then LUB lies between  $\ln 2$  and  $2(\sqrt{2} - 1)$ .

### 7.2 Speedup Factors

It has been shown [43, 45, 51, 53, 228] that the mixed criticality schedulability problem (preemptive or non-preemptive) is strongly NP-hard even if there are only two criticality levels. Hence only sufficient rather than exact analysis is possible. A list of open problems with regard to the schedulability of MCS was provided by Ekberg and Yi [168] in 2015. More recently (in 2017) Kahil et al. [271, 272] claim to have found a counter-example to the proof that the mixed-criticality optimality problem belongs to the class NP.

For approaches and tests that are only sufficient, an assessment of their quality is possible if a *speedup factor* can be computed. A speedup factor [275] of  $X$  ( $X > 1$ ) for schedulability test  $S$  implies that a task set that is schedulable on a processor of speed 1 will be deemed schedulable by  $S$  if the processor's speed is increased to  $X$ . Of course, in general, it is not possible to know if the task set is schedulable on the original speed 1 processor (this would require an exact and possibly even

a clairvoyant test), but a real scheduling scheme and test with a speedup factor of say 2 is clearly better than one with a speed up factor of 10.

For job-based fixed priority scheduling, a priority assignment scheme and test has been found [52, 53, 73, 307] with a speed up factor of  $S_L$  (for  $L$  criticality levels), where  $S_L$  is the root of the equation  $x^L = (1 + x)^{L-1}$ . For  $L = 2$  the result is  $S_2 = (1 + \sqrt{5})/2$  which is equal to the golden ratio,  $\phi = 1.618$ . This can be compared with a partitioned approach (all *HI*-criticality jobs have priorities higher than all *LO*-criticality jobs) which has an unbounded speedup factor. This latter result is easily illustrated by considering a two job system. The *LO*-criticality job has a small computation time, 1, and deadline of 2. The *HI*-criticality job has a huge computation time of  $G$  and a deadline of  $G + 1$ . These two jobs will both meet their deadlines if the *LO*-criticality task is given the highest priority. But the reverse priority assignment (which executes the *HI*-criticality job first) will only be schedulable if  $G + 1$  will fit into the deadline 2. To obtain this a speedup of  $(G + 1)/2$  is required. As  $G$  can be arbitrary large the speedup factor is effectively unbounded.

For EDF scheduled systems Baruah et al. [56, 73] prove that a variant of EDF (EDF-VD, described in Section 3.3) in which *HI*-criticality sporadic tasks (in a dual-criticality system) have their deadlines reduced (in the *LO*-criticality mode) is also schedulable on a single processor that is speeded up by a factor  $\phi$ . They also show that a finite set of independent jobs? scheduled on  $m$  identical multiprocessors is schedulable with a speed-up factor of  $\phi + 1 + 1/m$ . And on a partitioned system a speed-up factor of  $\phi + \epsilon$  is derivable for any value of  $\epsilon > 0$ . In later work [55] they improve this bound to  $4/3$  (1.333) rather than  $\phi$  (1.618). Further formal analysis of EDF-VD is provided by Li [305], Muller and Masrur [345] and Gu and Easwaran [211]. The MC-Fluid approach also has a speed-up factor of  $4/3$  [66], as does EDF-VD when applied to systems with degraded/imprecise guarantees [321].

### 7.3 Formal Language and Modelling Issues

The application of formal design languages, such as real-time BIT (Behavior Interactions Priorities), are being used to model MCS. And verification approaches such as model checking or simulation are being applied to both application software and multi-core platforms. See, for example, the work of Soggi et al. [418]. State-space explosion is of course always an issue with these approaches.

A 2005 paper by Amey et al. [19], which predates Vestal's work, looked at the (smart) certification of mixed criticality systems. They report real industrial application of formal code analysis to prove isolation between tasks of difference criticality levels. In one application, concerning safety-critical landing guidance

for ship-borne helicopters, SIL 4 code (the highest in UK Defence Standard 00-55 [340]) was executing in the same processor and memory space as SIL3 and SIL2 code. Another application (a civil jet engine monitoring unit) have Level-C and Level-E code co-located (DO-178B standard [395] has levels A down to E). In both of these examples formal analysis of information flow at the program level was able to demonstrate code segregation.

Compile time checking is also advocated by Lindgren et al. [315] with their experimental RTFM-language. Language constructs allow static assessment of the interfaces between critical and non-critical code. At run-time however separation is achieved by assigning higher priorities to the critical tasks. As indicated earlier this is not a very effective strategy in terms of efficient resource usage.

Model-based design using Synchronous Reactive models is used to design embedded control systems and is formalised within languages such as LUSTRE, SIGNAL and Simulink. Its application to MCS is considered in detail by Zhao et al [474]; where they adopt the elastic mixed-criticality task model for fixed-priority scheduling.

An alternative task model to that which underpins most scheduling research on MCS is the sporadic DAG (Directed Acyclic Graph) task model. An initial study, within the context of multiprocessor federated systems, is provided by Baruah [49]. Wang and Wang [459] and Medina et al. [336] also consider this model.

## 8 Systems Issues

A fundamental issue with MCS is separation. Many of the more theoretical papers reviewed here assume various levels of run-time monitoring and control. However, few papers consider or demonstrate how the required mechanisms can be implemented, Neukirchner et al. [353] presented one such paper. They consider memory protection, timing fault containment, admission control and (re-)configuration middleware for MCS. Their framework [179] is aimed at supporting AUTOSAR conforming applications within the automotive domain. An early paper looking at non-interference at the memory level for IMA platforms within the avionics industry is that of Hill and Lake [239].

Another detailed study of the overheads for two common implementations schemes for MCS is presented by Sigrist et al. [415]. They conclude that overheads of up to 97% can be encountered and they recommend that all scheduling models be extended to include parameters to capture the impact of run-time overheads. One source of overhead is the context switch time between tasks. Davis et al. [140] note that switching between tasks of the same criticality should be quicker than between tasks of different criticality. They produce analysis for the standard

fixed priority approach for MCS that explicitly caters for varying context switch times; they note however that priority assignment is not straightforward for this model [141].

The issue of monitoring is also addressed by Motruk et al. [343] in the context of their IDAMC (Integrated Dependable Architecture for Many Cores). This work builds on the more general (i.e. not MCS specific) work on separation, isolation and monitoring for SoC/NoC architectures. The MultiPARTES project is addressing virtualisation in terms of Model Driven Engineering for MCS [16, 127, 128, 447, 448]. Goossens et al. [197] are also looking at virtualisation “to allow independent design, verification and execution” with their CompSOC architecture. Paravirtualisation of legacy RTOSs to provide the necessary memory isolation is considered by Armbrust et al. [22].

Hypervisor technology is also being used to give the appropriate level of isolation in MCS. The DREAMS architecture uses it [294] to minimise interference via modelling patterns of execution. Evripidou and Burns [178] employ different execution-time servers (deferrable server for short deadline event-triggered work, and periodic server for periodic work) under the control of a hypervisor to bound the overheads associated with server technology. If there is a criticality induced mode change then the deferrable servers are transposed to the much more efficient (but less responsive) periodic servers. General hypervisor architectures for multi-core MCS are proposed by Pérez et al. [378] and Avramenko et al. [29]. Within the context of ARMv8-A, Lucas et al. [325] have developed a VOSYSmonitor that allows a safety-critical RTOS and a non-critical general purpose OS to co-exist on the same hardware platform. With a large Nuclear Fusion project, Cinque et al. [131] argue that virtualisation will not scale to the ‘tens of thousands’ of threads envisaged in this application. They propose a more lightweight solution using containers [130].

The issue of monitoring and testing MCS, including the use of HIL (hardware-in-the-loop), is addressed in a number of papers. Particular attention been given to multi-core platforms [176, 177, 362].

The development of purpose built hardware (FPGA based) to support reliable MCS is being undertaken as part of the RECOMP project. They aim to reduce the cost of certification for MCS on multiprocessor architectures by use of open source hardware and software [338, 354, 381]. Santos et al. [402] are also looking at systems built on FPGA platforms. They have developed a criticality-aware scrubbing mechanism that improves system reliability by up to 79%. Scrubbing is a technique for recovering from SEU (single event upsets) that effect FPGA platforms in harsh environments such as space.

A separation kernel has been developed by West et al. [311, 341, 462]. They can host guest operating systems, such as Linux or their own real-time operating



system (RTOS), QUEST-V. They partition the available cores into Sandboxes that have different criticality levels. Their architecture is aimed at achieving efficient resource partitioning and performance isolation. One means of achieving this is for interrupts to go directly to the appropriate partition, they do not have to be first handled by the hypervisor. An RTOS that addresses partitioning for the NoC (network-on-chip) as well as the processor is developed by Esposito et al. [175].

PikeOS [274] also employs a separation microkernel to provide ‘a powerful and efficient paravirtualization real-time operating system’ [397] for a partitioned multi-core platform. Vanga et al. [450] present a case study in the use of PikeOS in which they aim to give effective support to low criticality tasks that have short deadlines (low latency). Another high-assurance (micro) kernel is sel\_4. Lyons and Heiser [326] show how the sel\_4 model can be extended to cater for mixed criticality. VMs (virtual machines) that are appropriate for real-time Java-based mixed-criticality systems have been designed by Ziarek and Blanton [481] and Hamza et al. [225].

To implement the criticality mode change the run-time support system must support execution time monitoring, the modes and mode changes. Baruah and Burns [58] show how this can be achieved within the facilities provided by the Ada programming language. Kim and Jin do the same for a standard RTOS [281]. They make use of bitmaps to provide a very efficient implementation. DMPL [116] is a language designed specifically for distributed real-time MCS.

A further operating system designed to support mixed criticality is Kron-OS [138]. This controls the execution of RSFs (Repetitive Sequence of Frames) that is partitioned between two criticality levels. Instead of a purpose built RTOS a COTS platform is preferred by Raghenzani et al. [391]. They focus on Linux and attempt to characterise the interference that the platform can produce. HIPPEROS [366] is a multi-core OS designed for use in the avionics domain. They use one core to make all scheduling decisions. They also use the elastic task model (see Section 3.3) to allow *LO*-criticality tasks to degrade when an overrun occurs. Vetter et al. [452] show how a low latency criticality-aware network channel can be constructed between different OSs, in particular between Linux and one utilising ARM TrustZone.

As an alternative to using an RTOS to give the right level of protection and (safe) resources sharing, Zimmer et al. [482] have designed a processor (Flex-PRET) to directly support MCS. They use fine-grained multithreading and scratchpad memory to give protection to hard real-time tasks while increasing the resource utilisation of soft tasks. In effect soft tasks (threads) can safely exploit the spare capacity generated from the hard tasks at the cycle level. They have a soft-core FPGA implementation that caters for up to 8 hardware threads, each of which can support a number of software threads. A more focused scheme aimed at partitioning the

cache is described by Lesage et al. [303]. An LC (least critical) cache replacement policy is evaluated by Kumar et al. [289]. The effective use of cache, for a multi-core platform, is also considered by Chrisholm et al. [125]. Similarly, Ali Awan et al. incorporate criticality into their dynamic last-level cache partitioning scheme [33], and memory bandwidth regulation [30]. Partial lockdown and cache reclamation is also considered by this group [85]. Sritharan et al. [424] propose a time-based cache coherence protocol for bounding latencies from critical memory accesses.

A hardware platform that supports applications of different criticality must manage its I/O functions in a partitioned and hence safe (and secure) way. If lower criticality work can cause interrupt to occur ‘at any time’ then unpredictable overheads may be suffered by high criticality applications. This is a topic addressed by Paulitsch et al. [374]. They rightly claim that this topic is often overlooked.

Although research on MCS has generated many different approaches, there have been few empirical benchmarks or comparative studies. One useful study however was published in 2012 by Huang et al. [245]. They compared Vestal’s scheme with its optimal priority assignment, their improved slack scheduling scheme and Period Transformation (PT) (see Section 5.1). They conclude that Vestal’s approach and period transformation usually, though not always, outperform slack scheduling; and that there are additional overheads with period transformation and slack scheduling. Nevertheless the overheads were not excessive, typically an extra 0.3%. Later Fleming and Burns [182] compared Vestal’s approach, AMC (see Section 3.2) and PT for multiple criticality levels. As the number of criticality levels increased the relative advantage of PT, even when overheads are ignored, was observed to decrease. This observation was also supported by Huang et al. [246] who updated their study in 2014 and concluded that AMC-based scheduling gave the best performance for fixed priority sporadic task systems. This study also looked at the overheads involved in user-space implementation of AMC on top of Linux, without kernel modifications.

The need for useful benchmarks is noted in a number of papers. One industrially inspired case study is provided by Harbin et al. [229]. The use of realistic simulations to evaluate schemes is discussed by Bate et al. [77, 78], Griffin et al [205] and Ittersshagen et al. [260]. A brief comparison of approaches to multiprocessor scheduling of MCS is provided by Osmolovskiy et al. [365]. The evaluation of communications within a MCS is considered by the work of Napier et al. [349] and Petrakis et al. [380].

Another systems issue of crucial important in many mobile embedded systems is power consumption. The work of Broekaert et al. [91] allocates and monitors power budgets to different criticality levels. If a crucial VM (Virtual Machine) “*overpassed its power budget during its time partition, the extra power consumed*

*will be removed from the initial power budget of the next low critical VM scheduled*". Energy consumption is also addressed by Legout et al. [300]. They trade energy usage with deadline misses of low-criticality tasks, and claim a 17% reduction in energy with deadline misses kept below 4%. The objective of minimising energy usage is used by Zhang et al. [473] to drive task allocation in a multiprocessor system. As discussed earlier, a slightly different approach is taken by Huang et al. [250]. They advocate the use of DVFS (Dynamic Voltage and Frequency Scaling) to increase the speed of the processor if *HI*-criticality tasks need more than their  $C(LO)$  requirement. Hence *LO*-criticality work is not abandoned, but more energy is used. They integrate their approach with the EDF-VD scheduling scheme (see Section 3.3) and have, more recently, addressed multi-core platforms [350]. This approach is extended by Ali et al. [14] who propose a new dynamic power-aware scheduling scheme for hardware with discrete frequency levels. Awan et al. [34] consider how energy-aware task allocation can be utilised in the context of heterogeneous multi-core systems. A leakage aware DVFS scheme is proposed by Digalwar et al. [154] in 2017 and is shown to be more energy efficient when compared to existing approaches.

Where energy is limited or indeed the system is energy neutral, then criticality-aware energy usage becomes crucially important [454]. ENOS [458] is an experimental OS that addresses mixed resources (time and energy) and mixed criticality. It transforms the system through a series of 'energy modes' including one that ensures all state is safely stored in persistent memory before system blackout. Energy harvesting in the context of a battery-less real-time system is considered by Asyaban et al. [25]. They propose a scheduling scheme that satisfies both temporal and success-ratio constraints whilst addressing uncertainty in the platform's power management. Even where energy is not limited, isolation in terms of power usage and temperature control is important; an issue addressed by Grüttner in the context of heterogeneous MPSoCs [208]. Power management is also addressed by Juhasz and Jantsch [269] in order to manage QoS in a MCS.

Complex mixed criticality systems also present a number of significant challenges at the specification and design stage. Herrera et al. [236, 237] propose a modelling and design framework for MCS hosted on Systems-on-Chip and/or Systems-of-Systems. They present a core ontology but freely admit that there is considerable work to do before a sound engineering process is available for system builders/architects. Ittershagen et al. [261] go further and propose a systematic method for constructing workload and integration flow models for time-triggered MPSoC platforms. Design space exploration for MCS is considered by Muttillio et al. [346].

Giannopoulou et al. [190] support the development of MCS on multi-core platforms by the development of an appropriate tool chain. This group has also con-

sidered [193, 469] the mapping and design of fault-tolerant MCS to multi-core platforms.

## 9 Industry Practice and Standards

This survey covers the considerable body of research into MCS stemming from the model presented by Vestal [451]. Industry practice and safety standards; however, provide a somewhat different perspective on MCS [174]. There are different meanings assumed for some of the commonly used terms, and different objectives.

Determining the criticality of an application (or system function implemented via both hardware and software) is done via a system safety assessment that involves Hazard Analysis (HA) Failure Modes and Effects Analysis (FMAE) and Fault Trees Analysis (FTA). The criticality level typically depends on (i) an evaluation of the consequences of a failure, (ii) the probability that the failure occurs, and (iii) the provision of means to mitigate or cope with the fault. Hence the criticality level of an application may not necessary reflect the severity or consequences of failure. An example given by Esper et al. [173] and Ernst and Di Natale [172] comes from ISO 26262. If the probability of failure occurrence is very low, the ASIL level assigned may be low, despite severe consequences if a failure actually happens. A different application with a high probability of failure may be assigned a higher ASIL despite having lower severity consequences. With this interpretation, the idea of dropping low-criticality functionality in favour of completing that of higher criticality does not hold; the consequences would be more severe. ISO 26262 also permits high-criticality applications to be composed from lower criticality components with diverse implementations, again dropping one of the lower criticality components would remove the diversity and undermine the safety argument for the high-criticality function. The message here is that the criticality level is not the same as the importance of the application. Functionality that has low criticality cannot simply be dropped.

The standards require that “*sufficient independence*” or “*freedom from interference*” is demonstrated between functions of different criticality levels in both spatial and timing domains. If this is not done, then the whole system needs to be designed and developed according to methods appropriate for the highest criticality level involved, which would be untenable in practice for cost reasons. It remains a significant challenge to achieve the necessary separation, while also providing an efficient means of sharing resources. This is particularly apposite with the advent of multi-core and many-core platforms. Tool are required to verify that the necessary level of independence is indeed manifest in the design of MCS [453].

## 10 PhD Awards for Research within the Field of MCS

As an indication of the richness of the landscape for research within the many aspects of MCS we note (some) of the PhDs (or equivalent) that have been awarded to students studying topics within the domain of mixed criticality systems:

Mohemed El Mehdi Aichouch, *Evaluation of a Multiple-Criticality Real-Time Virtual Machine System and Configuration of an RTOS Resource Allocation Techniques*, INSA Rennes, 2014.

Christos Evripidou, *Scheduling for Mixed-criticality Hypervisor Systems in the Automotive Domain*, University of York, 2016.

Thomas Fleming, *Allocation and Optimisation of Mixed Criticality Cyclic Executives*, University of York, 2017.

Georgia Giannopoulou, *Implementation of Mixed-Criticality Applications on Multi-Core Architectures*, ETH Zurich, 2016.

Chetan Govindaiah, *Hardware Architecture Support for Mixed-Criticality and Real-Time Systems*, Iowa State University, 2016.

Romain Gratia, *A Generic, Efficient Approach to Mixed Criticality Sequencing for Multi-Core Processors*, Telecom Paristech, 2017.

Xiaozhe Gu, *Schedulability Analysis and Low-Criticality Execution Support for Mixed-Criticality Real-Time Systems on Uniprocessors*, Nanyang Technological University, 2017.

Zhishan Guo, *Real-Time Scheduling of Mixed-Criticality Workloads upon Platforms with Uncertainty*, University of North Carolina at Chapel Hill, 2016.

Biao Hu, *Schedulability Analysis of General Task Model and Demand Aware Scheduling in Mixed-Criticality Systems*, University of Munchen, 2017.

Pengcheng Huang, *Design and Optimisation of Mixed-Criticality Systems*, EFL Zurich, 2016.

Jaewoo Lee, *Resource-Efficient Scheduling of Multiprocessor Mixed-Criticality Real-Time Systems*, University of Pennsylvania, 2017.

Haohan Li, *Scheduling Mixed-Criticality Real-Time Systems*, University of North Carolina at Chapel Hill, 2013.

Dario Socci, *Scheduling of Certifiable Mixed-Criticality Systems*, University of Grenoble, 2016.

Jens Theis, *Certification-Cognizant Mixed-Criticality Scheduling in Time-Triggered Systems*, University of Kaiserslautern, 2015.

Iruno Agirre Troncoso, *Development and Certification of Mixed-Criticality Embedded Systems based on Probabilistic Timing Analysis*, Universitat Politècnica de Catalunya, 2018.

Hao Xu, *A Semi-Partitioned Model for Scheduling Mixed-Criticality Multi-Core Systems*, University of York, 2017.

Michael Zimmer, *Predictable Processors for Mixed-Criticality Systems and Precision-Timed I/O*, University of California at Berkeley, 2015.

## 11 Conclusion and Directions for Future Work

As identified in the introduction, the fundamental issue with MCS is how to reconcile the differing needs of separation (for safety) and sharing (for efficient resource usage). These concerns have led to somewhat of a bifurcation in the resulting research. Much of the implementation and systems work has concentrated on how to safely partition a system so that high-integrity components can, in some way, share computational and communication resources. By comparison, the more theoretical and scheduling research has largely focused on how criticality-specific worst-case execution times can be utilised to deliver systems that are schedulable at each criticality level but have high processor utilisation. Unfortunately these two areas of research are not easily integrated. Flexible scheduling requires, at least, dynamic partitioning. Certified systems require complete separation or at least static partitioning. Future work must address this mismatch.

A second topic for future work is a move away from a processor-centric view of MCS to one that incorporates other shared resources, for example communication; particularly on a multi-core or many-core platform. Can a shared bus provide the required separation, or is a Network-on-Chip protocol required? Work is only beginning to address these issues.

What becomes clear from reading the extensive literature that has been produced since the seminal paper of Vestal [451] in 2007, is that MCS presents a collection of interesting issues that are both theoretically intriguing and challenging from the perspective of implementation.

We finish this survey by listing open issues identified from reading the extensive research literature. (Many of these issues were presented by Alan Burns in his keynote talk at the Dagstuhl Seminar on Mixed Criticality Systems on Multi-core/Many-core Platforms in March 2015).

1. Holistic analysis is needed considering *all* system resources, particularly communications buses, networks, and access to memory, as well as the processor(s).
2. Appropriate models of system overheads and task dependencies are required, and need to be integrated into the analysis. In particular, attention needs to be paid to how overheads arising from tasks of one criticality level may impact tasks of different (particularly higher) criticality.
3. More work is needed to integrate run-time behaviour, i.e. monitoring and control, with the assumptions made during static analysis and verification.
4. Effective protocols are needed for sharing information between criticality levels.
5. There are a number of open issues with regards to graceful degradation and fault recovery. These include timely recovery back to the low-criticality mode of operation, and support for limited low-criticality functionality in higher criticality modes, avoiding the abandonment problem.
6. To be of practical use, techniques need to scale to more than two (possibly up to five) levels of criticality.
7. Better WCET analysis is needed to reduce the sound  $C(HI)$  and  $C(LO)$  estimates used, and to improve the confidence in these values.
8. How many different WCET estimates are required (or useful) for the same software?
9. Much of the existing research has looked at mixed criticality within a single scheduling scheme; however, further work is needed on integrating different schemes (e.g. cyclic executives for safety-critical applications, fixed priority for mission-critical applications, on the same processor).
10. Mechanisms are needed to tightly bound the impact of lower criticality tasks on those of higher criticality, independent of the behaviour or misbehaviour of the former, without significantly compromising performance, which may happen if strict isolation is enforced.
11. Time composability is needed across different criticality levels, so that the timing behaviour of tasks determined in isolation can be used when they are composed during system integration.

12. So far there has been little work on security as an aspect of criticality in real-time systems.
13. Probabilistic and statistical methods are a good match to requirements specified in terms of failure rates for different criticality levels; however, little work has been done on applying these techniques to MCS.
14. There are a number of formal aspects of scheduling still to be investigated.
15. Openly available benchmarks and case studies are needed for the evaluation of MCS techniques and analysis.
16. For research on MCS to have real impact it will be necessary to influence the relevant standards in the various application domains (e.g. automotive, aerospace).

Returning to the fundamental question underlying MCS research: how, in a disciplined way, to reconcile the conflicting requirements of *partitioning* for safety assurance and *sharing* for efficient resource usage. As yet we do not have the structures (models, methods, protocols, analysis etc.) needed to allow the tradeoffs between partitioning and separation to be properly evaluated. It is clear that MCS will continue to be a focus for practical and theoretical work for some time to come.

### **Acknowledgements**

The authors would like to thank Sanjoy Baruah for a number of very useful discussions on the topic of this paper. The research that went into writing this paper is funded in part by the ESPRC grants, MCC (EP/K011626/1) and MCCps (EP/P003664/1).



## References

- [1] L. Abdallah, M. Jan, J. Ermont, and C. Fraboul. I/O contention aware mapping of multi-criticalities real-time applications over many-core architectures. In *Proc. WiP, RTAS*, pages 25–28, 2016.
- [2] Y. Abdeddaim and D. Maxim. Probabilistic schedulability analysis for fixed priority mixed criticality real-time systems. In *Proc DATE*, pages 596–601, 2017.
- [3] A. Addisu, L. George, V. Sciandra, and M. Agueh. Mixed criticality scheduling applied to jpeg2000 video streaming over wireless multimedia sensor networks. In *Proc. WMC, RTSS*, pages 55–60, 2013.
- [4] K. Agrawal and S. Baruah. A Measurement-Based Model for Parallel Real-Time Tasks. In Sebastian Altmeyer, editor, *30th Euromicro Conference on Real-Time Systems (ECRTS 2018)*, volume 106 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 5:1–5:19. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018.
- [5] H. Ahmadian, F. Nekouei, and R. Obermaisser. Fault recovery and adaptation in time-triggered networks-on-chips for mixed-criticality systems. In *Proc. 12th International Symposium on Reconfigurable Communication-centric Systems-on-Chip (ReCoSoC)*, pages 1–8, 2017.
- [6] H. Ahmadian and R. Obermaisser. Time-triggered extension layer for on-chip network interfaces in mixed-criticality systems. In *Proc. Digital System Design (DSD)*, pages 693–699. IEEE, 2015.
- [7] Z. Al-Bayati, J. Caplan, B.H. Meyer, and H. Zeng. A four-mode model for efficient fault-tolerant mixed-criticality systems. In *Proc. DATE*, pages 97–102. IEEE, 2016.
- [8] Z. Al-Bayati, Q. Zhao, A. Youssef, H. Zeng, and Z. Gu. Enhanced partitioned scheduling of mixed-criticality systems on multicore platforms. In *20th Asia and South Pacific Design Automation Conference (ASP-DAC)*, pages 630–635, 2015.
- [9] B. Alahmad and S. Gopalakrishnan. A Risk-Constrained Markov Decision Process Approach to Scheduling Mixed-Criticality Job Sets. In *Proc 4th WMC (RTSS)*, 2016.
- [10] B. Alahmad and S. Gopalakrishnan. Isochronous execution models for mixed-criticality systems on parallel processors. In *WiP, RTSS*, pages 354–356, 2017.
- [11] B. Alahmad and S. Gopalakrishnan. Risk-aware scheduling of dual criticality job systems using demand distributions. *Leibniz Transactions on Embedded Systems*, 5(1):01–1–01:30, 2018.
- [12] B. Alahmad, S. Gopalakrishnan, L. Santinelli, and L. Cucu-Grosjean. Probabilities for mixed-criticality problems: Bridging the uncertainty gap. In *WiP, RTSS*, pages 1–4, 2011.
- [13] A. Ali and K.H. Kim. Cluster-based multicore real-time mixed-criticality scheduling. *Journal of Systems Architecture*, 79:45 – 58, 2017.
- [14] I. Ali, J. Seo, and K.H. Kim. A dynamic power-aware scheduling of mixed-criticality real-time systems. In *Computer and Information Technology; Ubiquitous Computing and Communications; Dependable, Autonomic and Secure Computing; Pervasive Intelligence and Computing (CIT/IUCC/DASC/PICOM)*, pages 438–445, 2015.
- [15] A. Alonso, J.A. de la Puente, J. Zamorano, M.A. de Miguel, E. Salazar, and J. Garrido. Safety concept for a mixed criticality on-board software system. *IFAC-PapersOnLine*, 48(10):240–245, 2015.

- [16] A. Alonso, C. Jouvray, S. Trujillo, M.A. de Miguel, C. Grepet, and J. Simo. Towards model-driven engineering for mixed-criticality systems: MultiPARTES approach. In *Proceedings of the Conference on Design, Automation and Test in Europe, WICERT, DATE*, 2013.
- [17] A. Alonso, E. Salazar, and M.A. de Miguel. A toolset for the development of mixed-criticality partitioned systems. In *HiPEAC Workshop*, 2014.
- [18] S. Altmeyer, L. Cucu-Grosjean, and R.I. Davis. Static probabilistic timing analysis for real-time systems using random replacement caches. *Real-Time Systems*, 51(1):77–123, 2015.
- [19] P. Amey, R. Chapman, and N. White. Smart certification of mixed criticality systems. In *Reliable Software Technologies, Proc. of the Ada Europe Conference*, pages 144–155. Springer Verlag, LNCS 3555, 2005.
- [20] J.H. Anderson, S.K. Baruah, and B.B. Brandenburg. Multicore operating-system support for mixed criticality. In *Proc. of the Workshop on Mixed Criticality: Roadmap to Evolving UAV Certification, San Francisco*, 2009.
- [21] R. Arbaud, D. Juhsz, and A. Jantsch. Resource management for mixed-criticality systems on multi-core platforms with focus on communication. Technical report, ResearchGate, 2018.
- [22] E. Armbrust, J. Song, G. Bloom, and G. Parmer. On spatial isolation for mixed-criticality, embedded systems. In L. Cucu-Grosjean and R. Davis, editors, *Proc. 2nd Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 15–20, 2014.
- [23] S. Asyaban and M. Kargahi. An exact schedulability test for fixed-priority preemptive mixed-criticality real-time systems. *Real-Time Systems Journal*, 54:32–90, 2018.
- [24] S. Asyaban and M. Kargahi. Feasibility interval for fixed-priority scheduling of mixed-criticality periodic tasks with offsets. *IEEE Embedded Systems Letters - online*, pages 1–4, 2018.
- [25] S. Asyaban, M. Kargahi, L. Thiele, and M. Mohaqeqi. Analysis and scheduling of a battery-less mixed-criticality system with energy uncertainty. *ACM Transactions on Embedded Computing Systems (TECS)*, 16(1):23, 2016.
- [26] N.C. Audsley. On priority assignment in fixed priority scheduling. *Information Processing Letters*, 79(1):39–44, 2001.
- [27] N.C. Audsley. Memory architectures for NoC-based real-time mixed criticality systems. In *Proc. WMC, RTSS*, pages 37–42, 2013.
- [28] N.C. Audsley, A. Burns, M. Richardson, K. Tindell, and A.J. Wellings. Applying new scheduling theory to static priority preemptive scheduling. *Software Engineering Journal*, 8(5):284–292, 1993.
- [29] S. Avramenko, S. Esposito, M. Violante, M. Sozzi, M. Traversone, M. Binello, and M. Terone. An hybrid architecture for consolidating mixed criticality applications on multicore systems. In *Proc. IEEE 21st International On-Line Testing Symposium (IOLTS)*, pages 26–29, 2015.
- [30] M.A. Awan, K. Bletsas, P. Souto, B. Akesson, and E. Tovar. Mixed-criticality scheduling with dynamic memory bandwidth. Technical Report TR-180603, CISTER/ISEP, 2018.
- [31] M.A. Awan, K. Bletsas, P. Souto, B. Akesson, E. Tovar, and J. Ali. Mixed-criticality scheduling with memory regulation. In *Proc. WiP, ECRTS*, page 22, 2016.
- [32] M.A. Awan, K. Bletsas, P. Souto, and E. Tovar. Semi-partitioned mixed-criticality scheduling. Technical report, CISTER/ISEP, 2016.

- [33] M.A. Awan, K. Bletsas, P.F. Souto, B. Akesson, and E. Tovar. Mixed-criticality scheduling with dynamic redistribution of shared cache. In Marko Bertogna, editor, *Proc. Euromicro Conference on Real-Time Systems (ECRTS)*, volume 76 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 18:1–21:20. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017.
- [34] M.A. Awan, D. Masson, and E. Tovar. Energy-aware task allocation onto unrelated heterogeneous multicore platform for mixed criticality systems. In *WiP, RTSS*, 2015.
- [35] M.A. Awan, D. Masson, and E. Tovar. Energy efficient mapping of mixed criticality applications on unrelated heterogeneous multicore platforms. In *Proc. 11th IEEE Symposium on Industrial Embedded Systems (SIES)*, pages 1–10, 2016.
- [36] P. Axer, M. Sebastian, and R. Ernst. Reliability analysis for mpsocs with mixed-critical, hard real-time constraints. In *Proceedings of the seventh IEEE/ACM/IFIP international conference on Hardware/software codesign and system synthesis, CODES+ISSS '11*, pages 149–158. ACM, 2011.
- [37] A. Azim and S. Fischmeister. Efficient mode changes in multi-mode systems. In *Proc. Computer Design (ICCD)*, pages 592–599. IEEE, 2016.
- [38] H. Baeh and J. Lee. Incorporating security constraints into mixed-criticality real-time scheduling. *IEICE Trans. on Information and Systems*, E100-D(9):2068–2080, 2017.
- [39] T.P. Baker. A stack-based resource allocation policy for realtime processes. In *Proc. IEEE Real-Time Systems Symposium (RTSS)*, pages 191–200, 1990.
- [40] J. Barhorst, T. Belote, P. Binns, J. Hoffman, J. Paunicka, P. Sarathy, J. Scoredos, P. Stanfill, D. Stuart, and R. Urzi. White paper: A research agenda for mixed-criticality systems, April 2009. Available at [http://www.cse.wustl.edu/~cdgill/CPSWEEK09\\_MCAR](http://www.cse.wustl.edu/~cdgill/CPSWEEK09_MCAR).
- [41] S. K. Baruah, A. Burns, and Z. Guo. Scheduling mixed-criticality systems to guarantee some service under all non-erroneous behaviours. In *Proc. ECRTS*, pages 131–140, 2016.
- [42] S.K. Baruah. Optimal utilization bounds for fixed priority scheduling of periodic task systems on identical multiprocessors. *IEEE Transactions on Software Engineering*, 53(6), 2004.
- [43] S.K. Baruah. Mixed criticality schedulability analysis is highly intractable. Technical report, University of North Carolina at Chapel Hill, 2009.
- [44] S.K. Baruah. Certification-cognizant scheduling of tasks with pessimistic frequency specification. In *Proc. 7th IEEE International Symposium on Industrial Embedded Systems (SIES'12)*, pages 31–38, 2012.
- [45] S.K. Baruah. Semantic-preserving implementation of multirate mixed criticality synchronous programs. In *Proc. RTNS*, 2012.
- [46] S.K. Baruah. Implementing mixed-criticality synchronous reactive programs upon uniprocessor platforms. *Real-Time Systems Journal*, 49(6), 2013.
- [47] S.K. Baruah. Implementing mixed criticality synchronous reactive systems upon multiprocessor platforms. Technical report, University of North Carolina at Chapel Hill, 2013.
- [48] S.K. Baruah. Response-time analysis of mixed criticality systems with pessimistic frequency specification. Technical report, University of North Carolina at Chapel Hill, 2013.
- [49] S.K. Baruah. The federated scheduling of systems of mixed-criticality sporadic DAG tasks. In *Proc. Real-Time Systems Symposium (RTSS)*, pages 227–236. IEEE, 2016.

- [50] S.K. Baruah. Schedulability analysis of mixed-criticality systems with multiple frequency specifications. In *Proc. International Conference on Embedded Software (EMSOFT)*, page 24. ACM, 2016.
- [51] S.K. Baruah, V. Bonifaci, G. D’Angelo, H. Li, A. Marchetti-Spaccamela, N. Megow, and L. Stougie. Scheduling real-time mixed-criticality jobs. In P. Hlinený and A.ín Kucera, editors, *Proc. of the 35th International Symposium on the Mathematical Foundations of Computer Science*, volume 6281 of *Lecture Notes in Computer Science*, pages 90–101. Springer, 2010.
- [52] S.K. Baruah, V. Bonifaci, G. D’Angelo, H. Li, A. Marchetti-Spaccamela, N. Megow, and L. Stougie. Mixed-criticality scheduling. In *10th Workshop on Models and Algorithms for Planning and Scheduling Problems (MAPSP)*, Nymburk, Czech Republic, pages 1–3, 2011.
- [53] S.K. Baruah, V. Bonifaci, G. D’Angelo, H. Li, A. Marchetti-Spaccamela, N. Megow, and L. Stougie. Scheduling real-time mixed-criticality jobs. *IEEE Transactions on Computers*, 61(8):1140–1152, 2012.
- [54] S.K. Baruah, V. Bonifaci, G. D’angelo, H. Li, A. Marchetti-Spaccamela, S. Van Der Ster, and L. Stougie. Preemptive uniprocessor scheduling of mixed-criticality sporadic task systems. *Journal of the ACM (JACM)*, 62(2):14, 2015.
- [55] S.K. Baruah, V. Bonifaci, G. D’Angelo, H. Li, A. Marchetti-Spaccamela, S. van der Ster, and L. Stougie. The preemptive uniprocessor scheduling of mixed-criticality implicit-deadline sporadic task systems. In *Proc. of ECRTS*, pages 145–154, 2012.
- [56] S.K. Baruah, V. Bonifaci, G. D’Angelo, A. Marchetti-Spaccamela, S. van der Ster, and L. Stougie. Mixed-criticality scheduling of sporadic task systems. In *Proc. of the 19th Annual European Symposium on Algorithms (ESA 2011) LNCS 6942, Saarbruecken, Germany*, pages 555–566, 2011.
- [57] S.K. Baruah and A. Burns. Sustainable schedulability analysis. In *Proc. of IEEE Real-Time Systems Symposium (RTSS)*, pages 159–168, 2006.
- [58] S.K. Baruah and A. Burns. Implementing mixed criticality systems in Ada. In A. Romanovsky, editor, *Proc. of Reliable Software Technologies - Ada-Europe 2011*, pages 174–188. Springer, 2011.
- [59] S.K. Baruah and A. Burns. Fixed-priority scheduling of dual-criticality systems. In *Proc. 21st RTNS*, pages 173–182. ACM, 2013.
- [60] S.K. Baruah and A. Burns. Achieving temporal isolation in multiprocessor mixed-criticality systems. In L. Cucu-Grosjean and R. Davis, editors, *Proc. 2nd Workshop on Mixed Criticality Systems (WMC)*, *RTSS*, pages 21–26, 2014.
- [61] S.K. Baruah, A. Burns, and R. I. Davis. Response-time analysis for mixed criticality systems. In *Proc. IEEE Real-Time Systems Symposium (RTSS)*, pages 34–43, 2011.
- [62] S.K. Baruah, A. Burns, and R.I. Davis. An extended fixed priority scheme for mixed criticality systems. In L. George and G. Lipari, editors, *Proc. ReTiMiCS, RTCSA*, pages 18–24, 2013.
- [63] S.K. Baruah and B. Chattopadhyay. Response-time analysis of mixed criticality systems with pessimistic frequency specification. In *Proc. RTCSA*, 2013.
- [64] S.K. Baruah, B. Chattopadhyay, H. Li, and I. Shin. Mixed-criticality scheduling on multiprocessors. *Real-Time Systems Journal*, 50:142–177, 2014.
- [65] S.K. Baruah, N. Cohen, G. Plaxton, and D. Varvel. Proportionate progress: A notion of fairness in resource allocation. *Algorithmica*, 15(6):600–625, 1996.

- [66] S.K. Baruah, A. Easwaran, and Z. Guo. MC-Fluid: Simplified and optimally quantified. In *Proc. IEEE Real-Time Systems Symposium (RTSS)*, pages 327–337, 2015.
- [67] S.K. Baruah, A. Easwaran, and Z. Guo. Mixed-criticality scheduling to minimize makespan. In *LIPICs-Leibniz International Proceedings in Informatics*, volume 65, 2016.
- [68] S.K. Baruah and G. Fohler. Certification-cognizant time-triggered scheduling of mixed-criticality systems. In *Proc. of IEEE Real-time Systems Symposium 2011*, December 2011.
- [69] S.K. Baruah and Z. Guo. Mixed criticality scheduling upon unreliable processors. Technical report, University of North Carolina at Chapel Hill, 2013.
- [70] S.K. Baruah and Z. Guo. Mixed-criticality scheduling upon varying-speed processors. In *Proc. IEEE 34th Real-Time Systems Symposium*, pages 68–77, 2013.
- [71] S.K. Baruah and Z. Guo. Scheduling mixed-criticality implicit-deadline sporadic task systems upon a varying-speed processor. In *Proc. IEEE Real-Time Systems Symposium*, pages 31–400. IEEE, 2014.
- [72] S.K. Baruah and Z. Guo. Mixed-criticality job models: a comparison. In L. Cucu-Grosjean and R. Davis, editors, *Proc. 3rd Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 1–5, 2015.
- [73] S.K. Baruah, H. Li, and L. Stougie. Mixed-criticality scheduling: Improving resource-augmented results. In *Computers and Their Applications, ISCA*, pages 217–223, 2010.
- [74] S.K. Baruah, H. Li, and L. Stougie. Towards the design of certifiable mixed-criticality systems. In *Proc. of the IEEE Real-Time Technology and Applications Symposium (RTAS)*, pages 13–22. IEEE, 2010.
- [75] S.K. Baruah and S. Vestal. Schedulability analysis of sporadic tasks with multiple criticality specifications. In *ECRTS*, pages 147–155, 2008.
- [76] S.K. Barugh. Scheduling analysis for a general model of mixed-criticality recurrent real-time tasks. In *Proc. IEEE RTSS*, pages 25–34, 2016.
- [77] I. Bate, A. Burns, and R.I. Davis. A bailout protocol for mixed criticality systems. In *Proc. 27th ECRTS*, pages 259–268, 2015.
- [78] I. Bate, A. Burns, and R.I. Davis. An enhanced bailout protocol for mixed criticality embedded software. *IEEE Transactions on Software Engineering*, 43(4):298–320, 2016.
- [79] L. Behera and P. Bhaduri. Time-triggered scheduling for multiprocessor mixed-criticality systems. In Atul Negi, Raj Bhatnagar, and Laxmi Parida, editors, *Proc. Distributed Computing and Internet Technology*, pages 135–151. Springer International Publishing, 2018.
- [80] G. Bernat and A. Burns. New results on fixed priority aperiodic servers. In *Proc. 20th IEEE Real-Time Systems Symposium*, pages 68–78, 1999.
- [81] G. Bernat and A. Burns. Multiple servers and capacity sharing for implementing flexible scheduling. *Real-Time Systems Journal*, 22:49–75, 2002.
- [82] G. Bernat, A. Colin, and S.M. Petters. Wcet analysis of probabilistic hard real-time systems. In *23rd IEEE Real-Time Systems Symposium*, pages 279–288. IEEE, 2002.
- [83] K.J. Biba. Integrity considerations for secure computer systems. Mtr-3153, Mitre Corporation, 1977.
- [84] E. Bini, M. Di Natale, and G.C. Buttazzo. Sensitivity analysis for fixed-priority real-time systems. In *Proc. ECRTS*, pages 13–22, 2006.

- [85] K. Bletsas, M.A. Awan, P.F. Souto, B. Akesson, and E. Tovar. Mixed-criticality systems with partial lockdown and cache reclamation upon mode change. In *Proc. WiP at ECRTS*, pages 22–24, 2017.
- [86] K. Bletsas and S.M. Petters. Using NPS-F for mixed criticality systems. In *Proc. WiP, RTSS*, page 25, 2012.
- [87] K. Bletsas and S.M. Petters. Using NPS-F for mixed criticality multicore systems. Cister-tr-130303, CISTER, 2013.
- [88] A. Blin, C. Courtaud, J. Sopena, and G. Muller. Maximizing parallelism without exploding deadlines in a mixed-criticality embedded system. In *Proc. ECRTS*, pages 109–119, 2016.
- [89] M. Bommert. Schedule-aware distributed of parallel load in a mixed criticality environment. In *Proc. JRWRTC, RTNS*, pages 21–24, 2013.
- [90] B.B. Brandenburg. A synchronous IPC protocol for predictable access to shared resources in mixed-criticality systems. In *Proc. IEEE Real-Time Systems Symposium*, pages 196–206. IEEE, 2014.
- [91] F. Broekaert, A. Fritsch, L. Sa, and S. Tverdyshev. Towards power-efficient mixed-critical systems. In *Proc. of OSPERT 2013*, pages 30–35, 2013.
- [92] I. Broster and A. Burns. An analysable bus-guardian for event-triggered communication. In *Proc. of the 24th Real-time Systems Symposium*, pages 410–419, Cancun, Mexico, Dec 2003. Computer Society, IEEE.
- [93] A. Burns. Preemptive priority based scheduling: An appropriate engineering approach. In S.H. Son, editor, *Advances in Real-Time Systems*, pages 225–248. Prentice-Hall, 1994.
- [94] A. Burns. The application of the original priority ceiling protocol to mixed criticality systems. In L. George and G. Lipari, editors, *Proc. ReTiMiCS, RTCSA*, pages 7–11, 2013.
- [95] A. Burns. System mode changes - general and criticality-based. In L. Cucu-Grosjean and R. Davis, editors, *Proc. 2nd Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 3–8, 2014.
- [96] A. Burns. An augmented model for mixed criticality. In Davis Baruah, Cucu-Grosjean and Maiza, editors, *Mixed Criticality on Multicore/Manycore Platforms (Dagstuhl Seminar 15121)*, volume 5(3), pages 92–93. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2015.
- [97] A. Burns and S.K. Baruah. Timing faults and mixed criticality systems. In Jones and Lloyd, editors, *Dependable and Historic Computing*, volume LNCS 6875, pages 147–166. Springer, 2011.
- [98] A. Burns and S.K. Baruah. Towards a more practical model for mixed criticality systems. In *Proc. 1st Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 1–6, 2013.
- [99] A. Burns and S.K. Baruah. Semi-partitioned cyclic executives for mixed criticality systems. In L. Cucu-Grosjean and R. Davis, editors, *Proc. 3rd Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 36–41, 2015.
- [100] A. Burns, R. Davis, S. K. Baruah, and I. Bate. Robust mixed-criticality systems. *IEEE Transactions on Computers – online*, 2018.
- [101] A. Burns and R.I. Davis. Mixed criticality on controller area network. In *Proc. Euromicro Conference on Real-Time Systems (ECRTS)*, pages 125–134, 2013.
- [102] A. Burns and R.I. Davis. Adaptive mixed criticality scheduling with deferred preemption. In *Proc. IEEE Real-Time Systems Symposium*, pages 21–30. IEEE, 2014.

- [103] A. Burns and R.I. Davis. Mixed criticality systems: A review. Technical Report MCC-1(H), available at <http://www-users.cs.york.ac.uk/burns/review.pdf>, Department of Computer Science, University of York, 2016.
- [104] A. Burns and R.I. Davis. Response-time analysis for mixed-criticality systems with arbitrary deadlines. In Agrawal and Easwaran, editors, *Proc. 5th Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 13–18, 2017.
- [105] A. Burns and R.I. Davis. A survey of research into mixed criticality systems. *ACM Computer Surveys*, 50(6):1–37, 2017.
- [106] A. Burns, T. Fleming, and S.K. Baruah. Cyclic executives, multi-core platforms and mixed-criticality applications. In *Proc. 27th ECRTS*, pages 3–12, 2015.
- [107] A. Burns, J. Harbin, and L.S. Indrusiak. A Wormhole NoC protocol for mixed criticality systems. In *Proc. IEEE Real-Time Systems Symposium*, pages 184–195. IEEE, 2014.
- [108] A. Burns and T.J. Quiggle. Effective use of abort in programming mode changes. *Ada Letters*, 1990.
- [109] A. Burns and A. J. Wellings. *Real-Time Systems and Programming Languages*. Addison Wesley Longman, 4th edition, 2009.
- [110] A. Burns and A.J. Wellings. Dual priority scheduling in Ada 95 and real-time POSIX. In *Proc. of 21th IFAC/IFIP Workshop on Real-Time Programming (WRTP96)*, 1996.
- [111] G. Buttazzo, G. Lipari, and L. Abeni. Elastic task model for adaptive rate control. In *IEEE Real-Time Systems Symposium*, pages 286–295, 1998.
- [112] G.C. Buttazzo. *Hard Real-Time Computing Systems*. Springer, 2005.
- [113] J. Caplan, Z. Al-bayati, H. Zeng, and B.H. Meyer. Mapping and scheduling mixed-criticality systems with on-demand redundancy. *IEEE Transactions on Computers*, PP(99):1–1, 2017.
- [114] G. Carvajal and S. Fischmeister. An open platform for mixed-criticality real-time ethernet. In *Proceedings of the Conference on Design, Automation and Test in Europe*, Proc. DATE, pages 153–156, 2013.
- [115] F.J. Cazorla, E. Quiones, T. Vardanega, L. Cucu-Grosjean, B. Triquet, G. Bernat, E.D. Berger, J. Abella, F. Wartel, M. Houston, L. Santinelli, L. Kosmidis, C. Lo, and D. Maxim. PROAR-TIS: Probabilistically analyzable real-time systems. *ACM Trans. Embedded Comput. Syst.*, 12(2):94, 2013.
- [116] S. Chaki and D. Kyle. Dmpl: Programming and verifying distributed mixed-synchrony and mixed-critical software. *Technical Report CMU/SEI-2016-TR-005*, 2016.
- [117] F. Checconi, T. Cucinotta, D. Faggioli, and G. Lipari. Hierarchical multiprocessor cpu reservations for the linux kernel. In *Proc. of 5th International Workshop on Operating Systems Platforms for Embedded Real-Time Applications (OSPert 2009)*, 2009.
- [118] G. Chen, N. Guan, B. Hu, and W. Yi. Edf-vd scheduling of flexible mixed-criticality system with multiple-shot transitions. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2018.
- [119] Y. Chen, K.G. Shin, and H. Xiong. Generalizing fixed-priority scheduling for better schedulability in mixed-criticality systems. *Information Processing Letters*, 116(8):508–512, 2016.
- [120] L. Cheng, K. Huang, G. Chen, B. Hu, and A. Knoll. Mixed-criticality control system with performance and robustness guarantees. In *Proc. IEEE Trustcom/BigDataSE/ICSS*, pages 767–775, 2017.

- [121] H. Chetto and M. Chetto. Some results of the earliest deadline scheduling algorithm. *IEEE Transactions on Software Engineering*, 15(10):1261–1269, 1989.
- [122] M. Chisholm, N. Kim, S. Tang, N. Otterness, J.H. Anderson, F.D. Smith, and D. Porter. Supporting mode changes while providing hardware isolation in mixed-criticality multicore systems. In *Proc. RTNS*, 2017.
- [123] M. Chisholm, N. Kim, B.C. Ward, N. Otterness, J.H. Anderson, and F.D. Smith. Reconciling the tension between hardware isolation and data sharing in mixed-criticality, multicore systems. In *Proc. Real-Time Systems Symposium (RTSS)*, pages 57–68. IEEE, 2016.
- [124] M. Chisholm, B. Ward, N. Kim, and J. Anderson. Cache sharing and isolation tradeoffs in multicore mixed-criticality systems. Technical report, University of North Carolina, 2015.
- [125] M. Chisholm, B. Ward, N. Kim, and J. Anderson. Cache-sharing and isolation tradeoffs in multicore mixed-criticality systems. In *Proc. IEEE Real-Time Systems Symposium (RTSS)*, pages 305–316, 2015.
- [126] H.S. Chwa, K.G. Shin, H. Baek, and J Lee. Physical-state-aware dynamic slack management for mixed criticality systems. In *Proc. RTAS*, 2018.
- [127] B. Cilku, A. Crespo, P. Puschner, J. Coronel, and S. Peiro. A memory arbitration scheme for mixed-criticality multicore platforms. In L. Cucu-Grosjean and R. Davis, editors, *Proc. 2nd Workshop on Mixed Criticality Systems (WMC)*, RTSS, pages 27–32, 2014.
- [128] B. Cilku, A. Crespo, P. Puschner, J. Coronel, and S. Peiro. A TDMA-based arbitration scheme for mixed-criticality multicore platforms. In *Proc EBCCSP*, pages 1–6. IEEE, 2015.
- [129] B. Cilku and P. Puschner. Towards temporal and spatial isolation in memory hierarchies for mixed-criticality systems with hypervisors. In L. George and G. Lipari, editors, *Proc. ReTiMiCS, RTCSA*, pages 25–28, 2013.
- [130] M. Cinque and D. Cotroneo. Towards lightweight temporal and fault isolation in mixed-criticality systems with real-time containers. In *Proc. 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W)*, pages 59–60, 2018.
- [131] M. Cinque and G.D. Tommasi. Real-time containers for large-scale mixed-criticality systems. In *WiP, RTSS*, pages 369–371, 2017.
- [132] A. Cohen, V. Perrelle, D. Potop-Butucaru, E. Soubiran, and Z. Zhang. Mixed-criticality in railway systems: A case study on signaling application. *Ada User Journal, Proc of Workshop on Mixed Criticality for Industrial Systems (WMCIS'2014)*, 35(2):138–143, 2014.
- [133] A. Crespo, A. Alonso, M. Marcos, J.A. Puente, and P. Balbastre. Mixed criticality in control systems. In *Proc. 19th World Congress The Federation of Automatic Control*, pages 12261–12271, 2014.
- [134] O. Cros, F. Fauberteau, L. George, and X. Li. Mixed-criticality over switched ethernet networks. *Ada User Journal, Proc of Workshop on Mixed Criticality for Industrial Systems (WMCIS'2014)*, 35(2):138–143, 2014.
- [135] O. Cros, L. George, and X.Li. A protocol for mixed-criticality management in switched ethernet networks. In L. Cucu-Grosjean and R. Davis, editors, *Proc. 3rd Workshop on Mixed Criticality Systems (WMC)*, RTSS, pages 12–17, 2015.
- [136] L. Cucu-Grosjean. Independence - a misunderstood property of and for probabilistic real-time systems. In N. Audsley and S.K. Baruah, editors, *In Real-Time Systems: the past, the present and the future*, pages 29–37, 2013.



- [137] L. Cucu-Grosjean, L. Santinelli, M. Houston, C. Lo, T. Vardanega, L. Kosmidis, J. Abella, E. Mezzetti, E. Quiones, and F.J. Cazorla. Measurement-based probabilistic timing analysis for multi-path programs. In *Proc. 24th Euromicro Conference on Real-Time Systems (ECRTS)*, pages 91–101, 2012.
- [138] V. David, A. Barbot, and D. Chabrol. Dependable real-time system and mixed criticality: Seeking safety, flexibility and efficiency with Kron-OS. *Ada User Journal*, 35(4):259–265, 2014.
- [139] R.I. Davis. *On exploiting spare capacity in hard real-time systems*. PhD thesis, University of York, UK, 1995.
- [140] R.I. Davis, S. Altmeyer, and A. Burns. Mixed criticality systems with varying context switch costs. In *Proc. RTAS*, 2018.
- [141] R.I. Davis, S. Altmeyer, and A. Burns. Priority assignment in fixed priority pre-emptive systems with varying context switch costs. In *Proc. RTSS Workshop on Open Problems*, pages 11–12, 2018.
- [142] R.I. Davis and M. Bertogna. Optimal fixed priority scheduling with deferred pre-emption. In *Proc. IEEE Real-Time Systems Symposium*, pages 39–50, 2012.
- [143] R.I. Davis and A. Burns. Hierarchical fixed priority preemptive scheduling. In *Proc. of IEEE Real-Time Systems Symposium (RTSS)*, pages 389–398, 2005.
- [144] R.I. Davis and A. Burns. Resource sharing in hierarchical fixed priority preemptive systems. In *Proceeding IEEE Real-Time Systems Symposium (RTSS)*, 2006.
- [145] R.I. Davis and A. Burns. Robust priority assignment for fixed priority real-time systems. In *Proc. of IEEE Real-Time Systems Symposium (RTSS)*, 2007.
- [146] R.I. Davis, A. Burns, R.J. Bril, and J.J. Lukkien. Controller area network (CAN) schedulability analysis: Refuted, revisited and revised. *Journal of Real-Time Systems*, 35(3):239–272, 2007.
- [147] R.I. Davis, L. Santinelli, S. Altmeyer, C. Maiza, and L. Cucu-Grosjean. Analysis of probabilistic cache related pre-emption delays. In *ECRTS*, pages 129–138, 2013.
- [148] R.I. Davis, K. Tindell, and A. Burns. Scheduling slack time in fixed priority preemptive systems. In *Proc. 14th IEEE Real-Time Systems Symposium*, 1993.
- [149] R.I. Davis and A. J. Wellings. Dual priority scheduling. In *Proc. 16th IEEE Real-Time Systems Symposium*, pages 100–109, 1995.
- [150] D. de Niz, B. Andersson, H. Kim, M. Klein, L.T.X. Phan, and R. Rajkumar. Mixed-criticality processing pipelines. In *Design, Automation Test in Europe Conference Exhibition (DATE)*, pages 1372–1375, 2017.
- [151] C. Deutschbein, T. Fleming, A. Burns, and S. Baruah. Multi-core cyclic executives for safety-critical systems. In Kim Guldstrand Larsen, Oleg Sokolsky, and Ji Wang, editors, *Proc. Dependable Software Engineering. Theories, Tools, and Applications, SETTA*, pages 94–109. Springer International Publishing, 2017.
- [152] J.L. Díaz, D.F. Garcia, K. Kim, C.G. Lee, L.L. Bello, J.M. López, and O. Mirabella. Stochastic analysis of periodic real-time systems. In *IEEE Real-Time Systems Symposium (RTSS)*, 2002.
- [153] J. Diemer and R. Ernst. Back suction: Service guarantees for latency-sensitive on-chip networks. In *Proceedings of the 2010 Fourth ACM/IEEE International Symposium on Networks-on-Chip*, Proc. NOCS '10, pages 155–162. IEEE Computer Society, 2010.

- [154] M. Digalwar, B.K. Raveendran, and S. Mohan. LAMCS: A leakage aware DVFS based mixed task set scheduler for multi-core processors. *Sustainable Computing: Informatics and Systems*, 15(Supplement C):63 – 81, 2017.
- [155] A.C. Dimopoulos, G. Bravos, G. Dimitrakopoulos, M. Nikolaidou, V. Nikolopoulos, and D. Anagnostopoulos. A multi-core context-aware management architecture for mixed-criticality smart building applications. In *Proc. System of Systems Engineering Conference (SoSE)*, pages 1–6. IEEE, 2016.
- [156] F. Dorin, P. Richard, M. Richard, and J. Goossens. Schedulability and sensitivity analysis of multiple criticality tasks with fixed-priorities. *Real-Time Systems Journal*, 46(3):305–331, 2010.
- [157] S. Draskovic, P. Huang, and L. Thiele. On the safety of mixed-criticality scheduling. In *Proc. 4th WMC (RTSS)*, page 6, 2016.
- [158] M. Dridi, M. Lallali, S. Rubin, F. Singhoff, and J-P. Diguët. Modeling and validation of a mixed-criticality noc router using the IF language. In *Proc. NoCArc, International Workshop on Network on Chip Architectures*. ACM, 2017.
- [159] M. Dridi, S. Rubin, M. Lallali, J. Sepulveda, F. Singhoff, and J-P. Diguët. DAS: An efficient noc router for mixed-criticality real-time systems. In *Proc. ICCD, Interenational Conference on Computer Design*. IEEE, 2017.
- [160] C. Dürr, Z. Hanzálek, C. Konrad, R. Sitters, O.C. Vásquez, and G. Woeginger. The triangle scheduling problem. *arXiv preprint arXiv:1602.04365*, 2016.
- [161] A. Easwaran. Demand-based scheduling of mixed-criticality sporadic tasks on one processor. In *Proc. IEEE 34th Real-Time Systems Symposium*, pages 78–87, 2013.
- [162] A. Easwaran and I. Shin. Compositional mixed-criticality scheduling. *CRTS 2014*, 2014.
- [163] L. Ecco, S. Tobuschat, S. Saidi, and R. Ernst. A mixed critical memory controller using bank privatization and fixed priority scheduling. In *Proc. Embedded and Real-Time Computing Systems and Applications (RTCSA)*, pages 1–10. IEEE, 2014.
- [164] S. Edgar and A. Burns. Statistical analysis of WCET for scheduling. In *Proc. 22nd IEEE Real-Time Systems Symposium*, 2001.
- [165] P. Ekberg, M. Stigge, N. Guan, and W. Yi. State-based mode switching with applications to mixed criticality systems. In *Proc. WMC, RTSS*, pages 61–66, 2013.
- [166] P. Ekberg and W. Yi. Bounding and shaping the demand of mixed-criticality sporadic task systems. In *ECRTS*, pages 135–144, 2012.
- [167] P. Ekberg and W. Yi. Bounding and shaping the demand of generalized mixed-criticality sporadic task systems. *Journal of Real-Time Systems*, 50:48–86, 2014.
- [168] P. Ekberg and W. Yi. A note on some open problems in mixed-criticality scheduling. In *Proc. RTOPS, 27th ECRTS*, pages 1–2, 2015.
- [169] P. Ekberg and W. Yi. Schedulability analysis of a graph-based task model for mixed-criticality systems. *Real-Time Systems*, 52:1–37, 2016.
- [170] P. Emberson and I. Bate. Minimising task migrations and priority changes in mode transitions. In *Proc. of the 13th IEEE Real-Time And Embedded Technology And Applications Symposium (RTAS 07)*, pages 158–167, 2007.
- [171] B. Engel. Tightening critical section bounds in mixed-criticality systems through preemptible hardware transactional memory. In *Proc. OSPERT*, pages 17–22, 2016.

- [172] R. Ernst and M. Di Natale. Mixed criticality systems? a history of misconceptions? *IEEE Design & Test*, 33(5):65–74, 2016.
- [173] A. Esper, G. Nelissen, V. Neils, and E. Tovar. How realistic is the mixed-criticality real-time system model. In *23rd International Conference on Real-Time Networks and Systems (RTNS 2015)*, pages 139–148, 2015.
- [174] A. Esper, G. Nelissen, V. Nélis, and E. Tovar. An industrial view on the common academic understanding of mixed-criticality systems. *Real-Time Systems*, 54(3):745–795, 2018.
- [175] S. Esposito, S. Avramenko, and M. Violante. RTOS for mixed criticality applications deployed on NoC-based COTS MPSoC. In *Proc. IEEE 19th Latin-American Test Symposium (LATS)*, pages 1–6, 2018.
- [176] S. Esposito, J. Sini, and M. Violante. Real-time validation of mixed-criticality applications. In *Proc. IEEE 19th Latin-American Test Symposium (LATS)*, pages 1–6, 2018.
- [177] S. Esposito, M. Violante, M. Sozzi, M. Terrone, and M. Traversone. A novel method for online detection of faults affecting execution-time in multicore-based systems. *ACM Trans. Embed. Comput. Syst.*, 16(4):94:1–94:19, 2017.
- [178] C. Evripidou and A. Burns. Scheduling for mixed-criticality hypervisor systems in the automotive domain. In *Proc. 4th WMC (RTSS)*, page 6, 2016.
- [179] G. Farrall, C. Stellwag, J. Diemer, and R. Ernst. Hardware and software support for mixed-criticality multicore systems. In *Proceedings of the Conference on Design, Automation and Test in Europe, WICERT, DATE*, 2013.
- [180] A. Finzi, A. MIFDAOUI, F. Frances, and E. Lochin. Mixed-Criticality on the AFDX Network: Challenges and Potential Solutions. In *Proc. 9th Embedded Real-Time Software and Systems (ERTS)*, pages pp. 1–9, 2018.
- [181] T. Fleming, S.K. Baruah, and A. Burns. Improving the schedulability of mixed criticality cyclic executives via limited task splitting. In *Proc. of the 24th International Conference RTNS*, pages 277–286. ACM, 2016.
- [182] T. Fleming and A. Burns. Extending mixed criticality scheduling. In *Proc. 1st Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 7–12, 2013.
- [183] T. Fleming and A. Burns. Incorporating the notion of importance into mixed criticality systems. In L. Cucu-Grosjean and R. Davis, editors, *Proc. 2nd Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 33–38, 2014.
- [184] T. Fleming and A. Burns. Investigating mixed criticality cyclic executive schedule generation. In L. Cucu-Grosjean and R. Davis, editors, *Proc. 3rd Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 42–47, 2015.
- [185] T. Fleming and A. Burns. Utilising asymmetric parallelism in multi-core mcs implemented via cyclic executives. In *Proc. 4th WMC (RTSS)*, page 6, 2016.
- [186] T. Fleming, H-M.Huang, A. Burns, C. Gill, S. Baruah, and C. Lu. Corrections to and discussion of "implementation and evaluation of mixed-criticality scheduling approaches for sporadic tasks". *ACM Trans. Embed. Comput. Syst.*, 16(3):77:1–77:4, 2017.
- [187] J. Freitag, S. Uhrig, and T. Ungerer. Virtual Timing Isolation for Mixed-Criticality Systems. In Sebastian Altmeyer, editor, *30th Euromicro Conference on Real-Time Systems (ECRTS 2018)*, volume 106 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 13:1–13:23. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018.

- [188] L. George, D. Masson, and V. Neli. Selective real-time data emission in mobile intelligent transport systems. In Agrawal and Easwaran, editors, *Proc. 5th Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 7–12, 2017.
- [189] O. Gettings, S. Quinton, and R.I. Davis. Mixed criticality systems with weakly-hard constraints. In *Proc. 23rd International Conference on Real-Time Networks and Systems (RTNS 2015)*, pages 237–246, 2015.
- [190] G. Giannopoulou and et al. DOL-BIP-Critical: A tool chain for rigorous design and implementation of mixed-criticality multi-core systems. Technical Report TR-2016-363, Verimag, 2016.
- [191] G. Giannopoulou, P. Huang, A. Gomez, and L. Thiele. Mixed-criticality runtime mechanisms and evaluation on multicore. In *Proc. RTAS*, 2015.
- [192] G. Giannopoulou, N. Stoimenov, P. Huang, and L. Thiele. Scheduling of mixed-criticality applications on resource-sharing multicore systems. In *Proc. Int. Conference on Embedded Software (EMSOFT)*, Montreal, 2013.
- [193] G. Giannopoulou, N. Stoimenov, P. Huang, and L. Thiele. Mapping mixed-criticality applications on multi-core architectures. In *Proc. Design, Automation and Test in Europe Conference and Exhibition (DATE)*, pages 1–6. IEEE, 2014.
- [194] G. Giannopoulou, N. Stoimenov, P. Huang, L. Thiele, and B. D. de Dinechin. Mixed-criticality scheduling on cluster-based manycores with shared communication and storage resources. *Real-Time Systems*, pages 1–51, 2015.
- [195] Robert Bosch GmbH. CAN specification version 2.0. Technical report, Postfach 30 02 40, D-70442 Stuttgart, 1991.
- [196] M. Gomony, J. Garside, B. Akesson, N. Audsley, and K. Goossens. A globally arbitrated memory tree for mixed-time-criticality systems. *IEEE Transactions on Computers*, 2016.
- [197] K. Goossens, A. Azevedo, K. Chandrasekar, M.D. Gomony, S. Goossens, M. Koedam, Y. Li, D. Mirzoyan, A. Molnos, A.B. Nejad, A. Nelson, and S. Sinha. Virtual execution platforms for mixed-time-criticality systems: The compsoc architecture and design flow. *SIGBED Rev.*, 10(3):23–34, 2013.
- [198] S. Goossens, B. Akesson, and K. Goossens. Conservative open-page policy for mixed time-criticality memory controllers. In *Proc. DATE*, pages 525–530, 2013.
- [199] S. Goossens, J. Kuijsten, B. Akesson, and K. Goossens. A reconfigurable real-time SDRAM controller for mixed time-criticality systems. In *Int'l Conf. on Hardware/Software Codesign and System Synthesis (CODES+ISSS)*, 2013.
- [200] D. Goswami, M. Lukasiewicz, R. Schneider, and S. Chakraborty. Time-triggered implementations of mixed-criticality automotive software. In *Proceedings of the Conference on Design, Automation and Test in Europe*, Proc. DATE, pages 1227–1232, 2012.
- [201] R. Gratia, T. Robert, and L. Pautet. Adaptation of RUN to mixed-criticality systems. In *Proc. 8th Junior Researcher Workshop on Real-Time Computing, RTNS*, 2014.
- [202] R. Gratia, T. Robert, and L. Pautet. Generalized mixed-criticality scheduling based on RUN. In *23rd International Conference on Real-Time Networks and Systems (RTNS 2015)*, pages 267–276, 2015.
- [203] R. Gratia, T. Robert, and L. Pautet. Scheduling of mixed-criticality systems with run. In *Proc. ETFA*, pages 1–8. IEEE, 2015.

- [204] P. Graydon and I. Bate. Safety assurance driven problem formulation for mixed-criticality scheduling. In *Proc. WMC, RTSS*, pages 19–24, 2013.
- [205] D. Griffin, I. Bate, B. Lesage, and F. Soboczenski. Evaluating mixed criticality scheduling algorithms with realistic workloads. In L. Cucu-Grosjean and R. Davis, editors, *Proc. 3rd Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 24–29, 2015.
- [206] S. Groesbrink, L. Almeida, M. de Sousa, and S.M. Petters. Towards certifiable adaptive reservations for hypervisor-based virtualization. In *Proc. of the 20th Real-Time and Embedded Technology and Applications Symposium (RTAS)*, 2014.
- [207] S. Groesbrink, S. Oberthir, and D. Baldin. Architecture for adaptive resource assignment to virtualized mixed-criticality real-time systems. In *Special Issue on the 4th Workshop on Adaptive and Reconfigurable Embedded Systems (APRES 2012)*, volume 10(1). ACM SIGBED Review, 2013.
- [208] K. Grüttner. Empowering mixed-criticality system engineers in the dark silicon era: Towards power and temperature analysis of heterogeneous mpsoCs at system level. In *Model-Implementation Fidelity in Cyber Physical System Design*, pages 57–90. Springer, 2017.
- [209] C. Gu, N. Guan, Q. Deng, and W. Yi. Partitioned mixed-criticality scheduling on multiprocessor platforms. In *Design, Automation and Test in Europe Conference and Exhibition (DATE), 2014*, pages 1–6. IEEE, 2014.
- [210] C. Gu, N. Guan, Q. Deng, and W. Yi. Improving ocbp-based scheduling for mixed-criticality sporadic task systems. In *Proc. RTCSA*, 2013.
- [211] X. Gu and A. Easwaran. Optimal speedup bound for 2-level mixed-criticality arbitrary deadline systems. In *Proceedings RTSOPS (ECRTS)*, pages 15–16, 2014.
- [212] X. Gu and A. Easwaran. Dynamic budget management with service guarantees for mixed-criticality systems. In *Proc. Real-Time Systems Symposium (RTSS)*, pages 47–56. IEEE, 2016.
- [213] X. Gu and A. Easwaran. Efficient schedulability test for dynamic-priority scheduling of mixed-criticality real-time systems. *ACM Trans. Embed. Comput. Syst.*, 17(1):24:1–24:24, 2017.
- [214] X. Gu, A. Easwaran, K.M. Phan, and I. Shin. Compositional mixed-criticality scheduling. Technical report, Nanyang Technological University, Singapore, 2014.
- [215] X. Gu, K.-M. Phan, A. Easwaran, and I. Shin. Resource efficient isolation mechanisms in mixed-criticality scheduling. In *Proc. 27th ECRTS*, pages 13–24. IEEE, 2015.
- [216] N. Guan, P. Ekberg, M. Stigge, and W. Yi. Effective and efficient scheduling of certifiable mixed-criticality sporadic task systems. In *IEEE RTSS*, pages 13–23, 2011.
- [217] N. Guan and W. Yi. Improving the scheduling of certifiable mixed criticality sporadic task systems. Technical report, University of Uppsala, 2012.
- [218] D. Guo and R. Pellizzoni. A requests bundling dram controller for mixed-criticality systems. In *Proc. IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 247–258, 2017.
- [219] Z. Guo. Mixed-criticality scheduling on varying-speed platforms with bounded performance drop rate. In *Proc. SMARTCOMP*, pages 1–3. IEEE, 2016.
- [220] Z. Guo and S. Baruah. Mixed-criticality real-time systems. In X. Wang, editor, *Cyber-Physical Systems: A Reference*. Springer, Berlin, Heidelberg, 2017.
- [221] Z. Guo and S.K. Baruah. Implementing mixed-criticality systems upon a preemptive varying-speed processor. *Leibniz Transactions on Embedded Systems*, 1(2):03–103:19, 2014.

- [222] Z. Guo and S.K. Baruah. The concurrent consideration of uncertainty in WCETs and processor speeds in mixed criticality systems. In *23rd International Conference on Real-Time Networks and Systems (RTNS 2015)*, pages 247–256, 2015.
- [223] Z. Guo, L. Santinelli, and K. Yang. EDF schedulability analysis on mixed-criticality systems with permitted failure probability. In *Proc. RTCSA*, 2015.
- [224] Z. Guo, S. Sruti, B. Ward, and S. Baruah. Sustainability in mixed-criticality scheduling. In *Proc. Real-Time Systems Symposium (RTSS)*, pages 24–33. IEEE, 2017.
- [225] H. Hamza, A. Hughes, and R. Kirner. On the design of a Java virtual machine for mixed-criticality systems. In *Proc. JTRES*. ACM, 2015.
- [226] J. J. Han, X. Tao, D. Zhu, and L. T. Yang. Resource sharing in multicore mixed-criticality systems: Utilization bound and blocking overhead. *IEEE Transactions on Parallel and Distributed Systems*, 28(12):3626–3641, 2017.
- [227] J.J. Han, X. Tao, D. Zhu, and H. Aydin. Criticality-aware partitioning for multicore mixed-criticality systems. In *Proc. Parallel Processing (ICPP)*, pages 227–235. IEEE, 2016.
- [228] Z. Hanzálek, T. Tunys, and P. Sucha. An analysis of the non-preemptive mixed-criticality match-up scheduling problem. *Journal of Scheduling*, pages 1–7, 2016.
- [229] J. Harbin, T. Fleming, L.S. Indrusiak, and A. Burns. GMCB: An industrial benchmark for use in real-time mixed-criticality networks-on-chip. In *Proc. WATERS, 27th ECRTS*, 2015.
- [230] P. Haririan and A. Garcia-Ortiz. A framework for hardware-based DVFS management in multicore mixed-criticality systems. In *Proc. 10th Symposium on Reconfigurable Communication-centric Systems-on-Chip (ReCoSoC)*, pages 1–7. IEEE, 2015.
- [231] M. Hassan and H. Patel. Criticality-and requirement-aware bus arbitration for multi-core mixed criticality systems. In *Proc. RTAS*, pages 1–11. IEEE, 2016.
- [232] M. Hassan, H. Patel, and R. Pellizzoni. A framework for scheduling dram memory accesses for multi-core mixed-time critical systems. In *Proc. RTAS*, pages 307–316. IEEE, 2015.
- [233] Mohamed Hassan. Heterogeneous mpsoCs for mixed criticality systems: Challenges and opportunities. *CoRR*, abs/1706.07429, 2017.
- [234] C. Herber, A. Richter, H. Rauchfuss, and A. Herkersdorf. Spatial and temporal isolation of virtual can controllers. In *Proc. ViRES, RTCSA*, 2013.
- [235] J. Herman, C. Kenna, M. Mollison, J. Anderson, and D. Johnson. RTOS support for multicore mixed-criticality systems. In *Proc. of the 18th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, 2012.
- [236] F. Herrera, S.H.A. Niaki, and I. Sander. Towards a modelling and design framework for mixed-criticality socs and systems-of-systems. In *Proc. 16th Euromicro Conf. on Digital Systems Design*, pages 989–996, 2013.
- [237] F. Herrera, P. Penil, and E. Villar. A model-based, single-source approach to design-space exploration and synthesis of mixed-criticality systems. In *Proc. SCOPES*, pages 88–91, 2015.
- [238] M. Hikmet, M.M. Kuo, P.S. Roop, and P. Ranjitkar. Mixed-criticality systems as a service for non-critical tasks. In *Proc. ISORC*, pages 221–228, 2016.
- [239] M.G. Hill and T.W. Lake. Non-interference analysis for mixed criticality code in avionics systems. In *Proceedings of the 15th IEEE international conference on Automated software engineering*, pages 257–260. IEEE Computer Society, 2000.

- [240] T. Hollstein, S.P Azad, T. Kogge, and B. Niazmand. Mixed-criticality NoC partitioning based on the NoCDepend dependability technique. In *Proc. 10th Symposium on Reconfigurable Communication-centric Systems-on-Chip (ReCoSoC)*, pages 1–8. IEEE, 2015.
- [241] P. Holman and J.H. Anderson. Adapting Pfair scheduling for symmetric multiprocessors. *Journal of Embedded Computing*, 1(4):543–564, 2005.
- [242] B. Hu, K. Huang, G. Chen, L. Cheng, and A. Knoll. Adaptive runtime shaping for mixed-criticality systems. In *Proc. 12th International Conference on Embedded Software, EMSOFT*, pages 11–20. IEEE Press, 2015.
- [243] B. Hu, K. Huang, G. Chen, L. Cheng, and A. Knoll. Adaptive workload management in mixed-criticality systems. *ACM Transactions on Embedded Computing Systems (TECS)*, 16(1):14, 2016.
- [244] B. Hu, K. Huang, P. Huang, L. Thiele, and A. Knoll. On-the-fly fast overrun budgeting for mixed-criticality systems. In *Proc. International Conference on Embedded Software (EMSOFT)*, pages 1–10. IEEE, 2016.
- [245] H-M. Huang, C. Gill, and C. Lu. Implementation and evaluation of mixed criticality scheduling approaches for periodic tasks. In *Proc. of the IEEE Real-Time Technology and Applications Symposium (RTAS)*, pages 23–32, 2012.
- [246] H-M. Huang, C. Gill, and C. Lu. Implementation and evaluation of mixed criticality scheduling approaches for sporadic tasks. *ACM Trans. Embedded Systems*, 13:126:1– 126:25, 2014.
- [247] P. Huang, G. Giannopoulou, R. Ahmed, D.B. Bartolini, and L. Thiele. An isolation scheduling model for multicores. In *Proc. IEEE Real-Time Systems Symposium (RTSS)*, pages 141–152, 2015.
- [248] P. Huang, G. Giannopoulou, N. Stoimenov, and L. Thiele. Service adaptations for mixed-criticality systems. Technical Report 350, ETH Zurich, Laboratory TIK, 2013.
- [249] P. Huang, G. Giannopoulou, N. Stoimenov, and L. Thiele. Service adaptations for mixed-criticality systems. In *Proc. 19th Asia and South Pacific Design Automation Conference (ASP-DAC)*, Singapore, 2014.
- [250] P. Huang, P. Kumar, G. Giannopoulou, and L. Thiele. Energy efficient DVFS scheduling for mixed-criticality systems. In *Proc. Embedded Software (EMSOFT)*, pages 1–10. IEEE, 2014.
- [251] P. Huang, P. Kumar, G. Giannopoulou, and L. Thiele. Run and be safe: mixed-criticality scheduling with temporal processor speedup. In *Proc. DATE*, 2015.
- [252] P. Huang, P. Kumar, N. Stoimenov, and L. Thiele. Interference constraint grapha new specification for mixed-criticality systems. In *Proc. 18th Emerging Technologies and Factory Automation (ETFA)*, pages 1–8. IEEE, 2013.
- [253] P. Huang, H. Yang, and L. Thiele. On the scheduling of fault-tolerant mixed-criticality systems. Technical report, Technical Report 351, ETH Zurich, Laboratory TIK, 2013.
- [254] P. Huang, H. Yang, and L. Thiele. On the scheduling of fault-tolerant mixed-criticality systems. In *Proc. Design Automation Conference (DAC)*, pages 1–6. IEEE, 2014.
- [255] B. Huber, C. El Salloum, and R. Obermaisser. A resource management framework for mixed-criticality embedded systems. In *34th IEEE IECON*, pages 2425–2431, 2008.
- [256] M. Hussein. Function allocation and bandwidth reservation for mixed-critical adaptive software systems. *Global Journal of Research in Engineering*, 18(4), 2018.

- [257] M. Hussein, A. Radermacher, and R. Nouacer. Model-based function mapping and bandwidth reservation for mixed-critical adaptive systems. In *Proc. Euromicro Conference on Digital System Design (DSD)*, pages 435–439, 2017.
- [258] L.S. Indrusiak, J. Harbin, and A. Burns. Average and worst-case latency improvements in mixed-criticality wormhole networks-on-chip. In *Proc. European/Euromicro Conference on Real-Time Systems (ECRTS)*, pages 47–56. IEEE, 2015.
- [259] S. Islam, R. Lindstrom, and N.Suri. Dependability driven integration of mixed criticality SW components. In *9th IEEE International Symposium on Object and Component-Oriented Real-Time Distributed Computing, ISORC 2006*, page 11, 2006.
- [260] P. Ittershagen, K. Gruttner, and W. Nebel. Mixed-criticality system modelling with dynamic execution mode switching. Technical report, Technical Report OFFIS, Oldenburg, Germany, 2015.
- [261] P. Ittershagen, K. Grüttner, and W. Nebel. An integration flow for mixed-critical embedded systems on a flexible time-triggered platform. *ACM Trans. Des. Autom. Electron. Syst.*, 23(4):51:1–51:25, 2018.
- [262] V. Izosimov and E. Levholt. Mixed criticality metric for safety-critical cyber-physical systems on multicore architectures. *MEDIAN: Methods*, 2(8), 2015.
- [263] J. Jalle, E. Quinones, J. Abella, L. Fossati, M. Zulianello, and F.J. Cazorla. A dual-criticality memory controller (DCmc): Proposal and evaluation of a space case study. In *Proc. IEEE Real-Time Systems Symposium*, pages 207–217. IEEE, 2014.
- [264] M. Jan, L. Zaourar, V. Legout, and L. Pautet. Handling criticality mode change in time-triggered systems through linear programming. *Ada User Journal, Proc of Workshop on Mixed Criticality for Industrial Systems (WMCIS'2014)*, 35(2):138–143, 2014.
- [265] M. Jan, L. Zaourar, and M. Pitel. Maximizing the execution rate of low criticality tasks in mixed criticality system. In *Proc. 1st WMC, RTSS*, pages 43–48, 2013.
- [266] X. Jin, J. Wang, and P. Zeng. End-to-end delay analysis for mixed-criticality WirelessHART networks. *Automatica Sinica*, 2(3):282–289, 2015.
- [267] X. Jin, C. Xia, H. Xu, J. Wang, and P. Zeng. Mixed criticality scheduling for industrial wireless sensor networks. *Sensors*, 16(9):1376, 2016.
- [268] C.B. Jones. Tentative steps toward a development method for interfering programs. *Transactions on Programming Languages and System*, 5(4):596–619, 1983.
- [269] D. Juhasz and A. Jantsch. Addressing the execution control problem in mixed-criticality systems. Technical report, ResearchGate, 2018.
- [270] N. Jung and J. Lee. *Improved Schedulability Analysis of Fixed-Priority for Mixed-Criticality Real-Time Multiprocessor Systems*, pages 1403–1409. Springer Singapore, Singapore, 2018.
- [271] R. Kahil, P. Poplavko, D. Socci, and S. Bensalem. Revisiting the computational complexity of mixed-critical scheduling. In Agrawal and Easwaran, editors, *Proc. 5th Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 25–30, 2017.
- [272] R. Kahil, P. Poplavko, D. Socci, and S. Bensalem. Revisiting the computational complexity of mixed-critical scheduling. Technical Report TR-2017-7, Verimag Research Report, 2017.
- [273] R. Kahil, P. Poplavko, D. Socci, and S. Bensalem. Predictability in mixed-criticality systems. Technical Report TR-2018-8, Verimag Research Report, 2018.
- [274] R. Kaiser. The PikeOS concept history and design,. Technical Report white paper, SYSGO, 2007.



- [275] B. Kalyanasundaram and K. Pruhs. Speed is as powerful as clairvoyance. *Journal of the ACM (JACM)*, 47(4):617–643, 2000.
- [276] C. Kamienski, M. Jentsch, M. Eisenhauer, J. Kiljander, E. Ferrera, P. Rosengren, J. Thestrup, E. Souto, W. S. Andrade, and D. Sadok. Application development for the internet of things: A context-aware mixed criticality systems development platform. *Computer Communications*, 2016.
- [277] O.R. Kelly, H. Aydin, and B. Zhao. On partitioned scheduling of fixed-priority mixed-criticality task sets. In *IEEE 10th International Conference on Trust, Security and Privacy in Computing and Communications*, pages 1051–1059, 2011.
- [278] H. Kim, D. Broman, E. Lee, M. Zimmer, A. Shrivastava, and J. Oh. A predictable and command-level priority-based DRAM controller for mixed-criticality systems. In *Proc. Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 317–326. IEEE, 2015.
- [279] N. Kim, B.C. Ward, M. Chisholm, J.H. Anderson, and F.D. Smith. Attacking the one-out-of-m multicore problem by combining hardware management with mixed-criticality provisioning. *Real-Time Systems*, 53(5):709–759, Sep 2017.
- [280] N. Kim, B.C. Ward, M. Chisholm, C-Y. Fu, J.H. Anderson, and F.D. Smith. Attacking the one-out-of-m multicore problem by combining hardware management with mixed-criticality provisioning. In *Proc. RTAS*, pages 1–12. IEEE, 2016.
- [281] Y.-S. Kim and H.-W. Jin. Towards a practical implementation of criticality mode change in RTOS. Technical report, Konkuk University, Korea, 2014.
- [282] A. Kostrzewa, S. Saidi, and R. Ernst. Dynamic control for mixed-criticality networks-on-chip. In *Proc. IEEE Real-Time Systems Symposium (RTSS)*, pages 317–326, 2015.
- [283] A. Kostrzewa, S. Tobuschat, L. Ecco, and R. Ernst. Adaptive load distribution in mixed-critical networks-on-chip. In *Proc. 22nd Asia and South Pacific Design Automation Conference (ASP-DAC)*, pages 732–737, 2017.
- [284] O. Kotaba, J. Nowotschy, M. Paulitschy, S.M. Petters, and H. Theiling. Multicore in real-time systems – temporal isolation challenges due to shared resources. In *Proc. Conference on Design, Automation and Test in Europe, WICERT, DATE*, 2013.
- [285] A. Kritikakou, O. Baldellon, C. Pagetti, C. Rochange, M. Roy, and F. Vargas. Monitoring on-line timing information to support mixed-critical workloads. In *WiP, RTSS*, pages 9–10, 2013.
- [286] A. Kritikakou, T. Marty, and M. Roy. DYNASCOPE: Dynamic software controller to increase resource utilization in mixed-critical systems. *ACM Trans. Des. Autom. Electron. Syst.*, 23(2):13:1–13:26, 2017.
- [287] A. Kritikakou, C. Pagetti, O. Baldellon, M. Roy, and C. Rochange. Run-time control to increase task parallelism in mixed-critical systems. In *ECRTS*, pages 119–128, 2014.
- [288] A. Kritikakou, C. Pagetti, C. Rochange, M. Roy, M. Faugre, S. Girbal, and D.G. Prez. Distributed run-time WCET controller for concurrent critical tasks in mixed-critical systems. In *Proc. RTNS*, 2014.
- [289] N.G. Kumar, S. Vyas, R.K. Cytron, C.D. Gill, J. Zambreno, and P.H. Jones. Cache design for mixed criticality real-time systems. In *Proc. ICCD*, pages 513–516. IEEE, 2014.
- [290] M. Kttler, M. Roitzsch, C-J. Hamann, and M. Vlp. Probabilistic analysis of low-criticality execution. In Agrawal and Easwaran, editors, *Proc. 5th Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 19–24, 2017.

- [291] A. Lackorzynski, A. Warg, M. Voelp, and H. Haertig. Flattening hierarchical scheduling. In *Proc. ACM EMSOFT*, pages 93–102, 2012.
- [292] K. Lakshmanan, D. de Niz, and R. Rajkumar. Mixed-criticality task synchronization in zero-slack scheduling. In *IEEE RTAS*, pages 47–56, 2011.
- [293] K. Lakshmanan, D. de Niz, R. Rajkumar, and G. Moreno. Resource allocation in distributed mixed-criticality cyber-physical systems. In *ICDCS*, pages 169–178, 2010.
- [294] A. Larrucea, I. Martinez, V. Brocal, S. Peirò, H. Ahmadian, J. Perez, and R. Obermaisser. DREAMS: Cross-domain mixed-criticality patterns. In *Proc. 4th WMC (RTSS)*, page 6, 2016.
- [295] J. Lee, H.S. Chwa, A. Easwaran, I. Shin, and I. Lee. Towards compositional mixed-criticality real-time scheduling in open systems. In L. Almeida and D. de Niz, editors, *Proc. 8th Workshop on Compositional Real-Time Systems (CRTS)*, RTSS, 2015.
- [296] J. Lee, H.S. Chwa, A. Easwaran, I. Shin, and I. Lee. Towards compositional mixed-criticality real-time scheduling in open systems: invited paper. *ACM SIGBED Review*, 13(3):49–51, 2016.
- [297] J. Lee, H.S. Chwa, L.T.X. Phan, I. Shin, and I. Lee. MC-ADAPT: Adaptive task dropping in mixed-criticality scheduling. *ACM Trans. Embed. Comput. Syst.*, 16:163:1–163:21, 2017.
- [298] J. Lee, K.-M. Phan, Z Gu, J. Lee, A. Easwaran, I. Shin, and I. Lee. MC-Fluid: Fluid model-based mixed-criticality scheduling on multiprocessors. In *Proc. IEEE Real-Time Systems Symposium*, pages 41–52. IEEE, 2014.
- [299] J. Lee, S. Ramanathan, K.-M. Phan, A. Easwaran, I. Shin, and I. Lee. Mc-fluid: Multi-core fluid-based mixed-criticality scheduling. *IEEE Transactions on Computers*, online, 2017.
- [300] V. Legout, M. Jan, and L. Pautet. Mixed-criticality multiprocessor real-time systems: Energy consumption vs deadline misses. In L. George and G. Lipari, editors, *Proc. ReTiMiCS, RTCSA*, pages 1–6, 2013.
- [301] J.P. Lehoczky and S. Ramos-Thuel. An optimal algorithm for scheduling soft-aperiodic tasks fixed-priority preemptive systems. In *Proc. 13th IEEE Real-Time Systems Symposium*, pages 110–123, 1992.
- [302] J.P. Lehoczky, L. Sha, and J.K. Strosnider. Enhanced aperiodic responsiveness in a hard real-time environment. In *Proc. 8th IEEE Real-Time Systems Symposium*, pages 261–270, 1987.
- [303] B. Lesage, I. Puaut, and A. Sez nec. PRETI: Partitioned real-time shared cache for mixed-criticality real-time systems. In *Proc. 20th RTNS*, pages 171–180, 2012.
- [304] J. Y-T. Leung and J. Whitehead. On the complexity of fixed-priority scheduling of periodic real-time tasks. *Performance Evaluation (Netherlands)*, 2(4):237–250, Dec. 1982.
- [305] H. Li. *Scheduling Mixed-Criticality Real-Time Systems*. PhD thesis, The University of North Carolina at Chapel Hill, 2013.
- [306] H. Li and S.K. Baruah. An algorithm for scheduling certifiable mixed-criticality sporadic task systems. In *Proc. of the Real-Time Systems Symposium*, pages 183–192, San Diego, CA, 2010. IEEE Computer Society Press.
- [307] H. Li and S.K. Baruah. Load-based schedulability analysis of certifiable mixed-criticality systems. In *Proc. EMSOFT*, pages 99–107. ACM Press, 2010.
- [308] H. Li and S.K. Baruah. Global mixed-criticality scheduling on multiprocessors. In *Proc. ECRTS*, pages 99–107. IEEE Computer Society Press, 2012.

- [309] J. Li, D. Ferry, S. Ahuja, K. Agrawal, C. Gill, and C. Lu. Mixed-criticality federated scheduling for parallel real-time tasks. In *Proc. RTAS*, pages 1–12. IEEE, 2016.
- [310] J. Li, D. Ferry, S. Ahuja, K. Agrawal, C. Gill, and C. Lu. Mixed-criticality federated scheduling for parallel real-time tasks. *Real-Time Systems*, 53(5):760–811, 2017.
- [311] Y. Li, R. West, and E. Missimer. The quest-v separation kernel for mixed criticality systems. In *Proc. 1st WMC, RTSS*, pages 31–36, 2013.
- [312] Z. Li and S. He. Fixed-priority scheduling for two-phase mixed-criticality systems. *ACM Trans. Embed. Comput. Syst.*, 17(2):35:1–35:20, 2018.
- [313] Z. Li and L. Wang. Memory-aware scheduling for mixed-criticality systems. In *Proc ICCSA*, pages 140–156. Springer, LNCS 9787, 2016.
- [314] J. Lin, A.M.K. Cheng, D. Steel, and M.Y.-C. Wu. Scheduling mixed-criticality real-time tasks with fault tolerance. In L. Cucu-Grosjean and R. Davis, editors, *Proc. 2nd Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 39–44, 2014.
- [315] P. Lindgren, D. Pereira, J. Eriksson, M. Lindner, and L. Miguel. Rtfm-lang static semantics for systems with mixed criticality. *Ada User Journal, Proc of Workshop on Mixed Criticality for Industrial Systems (WMCIS'2014)*, 35(2):128–132, 2014.
- [316] G. Lipari and S.K. Baruah. Greedy reclamation of unused bandwidth in constant bandwidth servers. In *IEEE Proc. of the 12th Euromicro Conference on Real-Time Systems*, Stockholm, Sweden, June 2000.
- [317] G. Lipari and E. Bini. A methodology for designing hierarchical scheduling systems. *J. Embedded Comput.*, 1(2):257–269, 2005.
- [318] G. Lipari and G. Buttazzo. Resource reservation for mixed criticality systems. In *Proceeding of Workshop on Real-Time Systems: the past, the present, and the future*, pages 60–74, York, UK, 2013.
- [319] C.L. Liu and J.W. Layland. Scheduling algorithms for multiprogramming in a hard real-time environment. *JACM*, 20(1):46–61, 1973.
- [320] D. Liu, N. Guan, J. Spasic, G. Chen, S. Liu, T. Stefanov, and W. Yi. Scheduling analysis of imprecise mixed-criticality real-time tasks. *IEEE Transactions on Computers*, 67(7):975–991, July 2018.
- [321] D. Liu, J. Spasic, N. Guan, G. Chen, S. Liu, T. Stefanov, and W. Yi. EDF-VD scheduling of mixed-criticality systems with degraded quality guarantees. In *Proc. IEEE RTSS*, pages 35–46, 2016.
- [322] G. Liu, Y. Lu, S. Wang, and Z. Gu. Partitioned multiprocessor scheduling of mixed-criticality parallel jobs. In *Proc. Embedded and Real-Time Computing Systems and Applications (RTCSA)*. IEEE, 2014.
- [323] J.W.S. Liu. *Real-Time Systems*. Prentice Hall, 2000.
- [324] J. López, J. Díaz, J. Entralgo, and D. García. Stochastic analysis of real-time systems under preemptive priority-driven scheduling. *Real-Time Systems*, pages 180–207, 2008.
- [325] P. Lucas, K. Chappuis, M. Paolino, N. Dagieu, and D. Raho. VOSYSmonitor, a low latency monitor layer for mixed-criticality systems on ARMv8-A. In Marko Bertogna, editor, *Proc. Euromicro Conference on Real-Time Systems (ECRTS)*, volume 76 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 6:1–6:18. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017.

- [326] A. Lyons and G. Heiser. Mixed-criticality support in a high-assurance, general-purpose micro-kernel. In L. Cucu-Grosjean and R. Davis, editors, *Proc. 2nd Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 9–14, 2014.
- [327] M. Mahdiani and A. Masrur. Introducing utilization caps into mixed-criticality scheduling. In *Proc. Digital System Design (DSD)*, pages 388–395. IEEE, 2016.
- [328] S.O. Marinescu, D. Tamas-Selicean, V. Acretoaie, and P. Pop. Timing analysis of mixed-criticality hard real-time applications implemented on distributed partitioned architectures. In *17th IEEE International Conference on Emerging Technologies and Factory Automation*, 2012.
- [329] A. Masrur. A probabilistic scheduling framework for mixed-criticality systems. In *Proc. DAC*, page 132. ACM, 2016.
- [330] A. Masrur, D. Muller, and M. Werner. Bi-level deadline scaling for admission control in mixed-criticality systems. In *Proc. 21st IEEE Embedded and Real-Time Computing Systems and Applications (RTCSA)*, pages 100–109. IEEE, 2015.
- [331] F.R. Massaro, E.L. Ursini, and P.O. Martins. Integrating proactive mode changes in mixed criticality systems. *arXiv [cs.OS]*, 2018. arXiv: 1806.11432.
- [332] S. Maurer and R. Kirner. Cross-criticality interfaces for cyber-physical systems. In *Proc. 1st IEEE Int'l Conference on Event-based Control, Communication, and Signal Processing*, 2015.
- [333] D. Maxim, R.I. Davis, L. Cucu-Grosjean, and A. Easwaran. Probabilistic analysis for mixed criticality scheduling with SMC and AMC. In *Proc. 4th WMC (RTSS)*, page 6, 2016.
- [334] D. Maxim, R.I. Davis, L. Cucu-Grosjean, and A. Easwaran. Probabilistic analysis for mixed criticality systems using fixed priority preemptive scheduling. In *Proc RTNS, International Conference on Real-Time Networks and Systems*, page 10, 2017.
- [335] R. Medina, E. Borde, and L. Pautet. Availability analysis for synchronous data-flow graphs in mixed-criticality systems. In *Proc. 11th IEEE Symposium on Industrial Embedded Systems (SIES)*, pages 1–6, 2016.
- [336] R. Medina, E. Borde, and L. Pautet. Directed acyclic graph scheduling for mixed-criticality systems. In Johann Blicberger and Markus Bader, editors, *Reliable Software Technologies – Ada-Europe*, pages 217–232. Springer International Publishing, 2017.
- [337] R. Medina, E. Borde, and L. Pautet. Availability enhancement and analysis for mixed-criticality systems on multi-core. In *Proc. Design, Automation Test in Europe Conference Exhibition (DATE)*, pages 1271–1276, 2018.
- [338] M. Mendez, J.L.G. Rivas, D.F. Garca-Valdecasas, and J. Diaz. Open platform for mixed-criticality applications. In *Proceedings of the Conference on Design, Automation and Test in Europe, WICERT, DATE*, 2013.
- [339] M-V. Micea, C-S. Stangaciu, V. Stangaciu, and D-I. Curiac. Novel hybrid scheduling technique for sensor nodes with mixed criticality tasks. *Sensors*, 17(7), 2017.
- [340] Ministry of Defence. Requirements for safety related software in defence equipment. Defence standard, 00-55, Ministry of Defence, 1997.
- [341] E. Missimer, K. Missimer, and R. West. Mixed-criticality scheduling with i/o. In *Proc. ECRTS*, pages 120–130, 2016.
- [342] M. Mollison, J. Erickson, J. Anderson, S.K. Baruah, and J. Scoredos. Mixed criticality real-time scheduling for multicore systems. In *Proc. of the 7th IEEE International Conference on Embedded Software and Systems*, pages 1864–1871, 2010.

- [343] B. Motruk, J. Diemer, R. Buchty, R. Ernst, and M. Berekovic. Idamc: A many-core platform with run-time monitoring for mixed-criticality. *Ninth IEEE International Symposium on High-Assurance Systems Engineering (HASE'05)*, pages 24–31, 2012.
- [344] S. Mubeen, M. Gålnander, J. Lundbäck, and K-L Lundbäck. Extracting timing models from component-based multi-criticality vehicular embedded systems. In Shahram Latifi, editor, *Information Technology - New Generations*, pages 709–718. Springer International Publishing, 2018.
- [345] D. Muller and A. Masrur. The scheduling region of two-level mixed-criticality systems based on EDF-VD. In *Proceedings of the Conference on Design, Automation and Test in Europe*, Proc. DATE, pages 978–981, 2014.
- [346] V. Muttillio, G. Valente, and L. Pomante. Criticality-aware design space exploration for mixed-criticality embedded systems. In *Companion of the ACM/SPEC International Conference on Performance Engineering*, ICPE '18, pages 45–46. ACM, 2018.
- [347] K. Nagalakshmi and N. Gomathi. Criticality-cognizant clustering-based task scheduling on multicore processors in the avionics domain. *International Journal of Computational Intelligence Systems*, 11:219–238, 2018.
- [348] M. Nager, M. Baunach, P. Priller, and J. Wurzinger. Real-time multiplexing of mixed-criticality data streams for automotive multi-core test systems. In *Proc. IEEE International Conference on Vehicular Electronics and Safety (ICVES)*, pages 220–227, 2017.
- [349] K. Napier, O. Horst, and C. Prehofer. Comparably evaluating communication performance within mixed-criticality systems. In *Proc. 4th WMC (RTSS)*, page 6, 2016.
- [350] S. Narayana, P. Huang, G. Giannopoulou, L. Thiele, and R.V. Prasad. Exploring energy saving for mixed-criticality systems on multi-cores. In *Proc. RTAS*, pages 1–12. IEEE, 2016.
- [351] M. Neukirchner, P. Axer, T. Michaels, and R. Ernst. Monitoring of workload arrival functions for mixed-criticality systems. In *Proc. IEEE 34th Real-Time Systems Symposium*, pages 88–96, 2013.
- [352] M. Neukirchner, S. Quinton, and K. Lampka. Multi-mode monitoring for mixed-criticality real-time systems. In *Int'l Conf. on Hardware/Software Codesign and System Synthesis (CODES+ISSS)*, 2013.
- [353] M. Neukirchner, S. Stein, H. Schrom, J. Schlatow, and R. Ernst. *Contract-based dynamic task management for mixed-criticality systems*, pages 18–27. IEEE, 2011.
- [354] R. Nevalainen, U. Kremer, O. Slotosch, D. Truscan, and V. Wong. Impact of multicore platforms in hardware and software certification. In *Proceedings of the Conference on Design, Automation and Test in Europe, WICERT, DATE*, 2013.
- [355] L.M. Ni and P.K. McKinley. A survey of wormhole routing techniques in direct networks. *Computer*, 26(2):62–76, Feb 1993.
- [356] D.de Niz, K. Lakshmanan, and R. Rajkumar. On the scheduling of mixed-criticality real-time task sets. In *Real-Time Systems Symposium*, pages 291–300. IEEE Computer Society, 2009.
- [357] D.de Niz and L.T.X. Phan. Partitioned scheduling of multi-modal mixed-criticality real-time systems on multiprocessor platforms. In *Proc. Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 111–122, April 2014.
- [358] D.de Niz, L. Wrage, A. Rowe, and R. Rajkumar. Utility-based resource overbooking for cyber-physical systems. In *Proc. RTCSA*, 2013.

- [359] A. Novak, P. Sucha, and Z. Hanzalek. Efficient algorithm for jitter minimization in time-triggered periodic mixed-criticality message scheduling problem. In *Proc. RTNS*, pages 23–31. ACM, 2016.
- [360] A. Novak, P. Sucha, and Z. Hanzalek. On solving non-preemptive mixed-criticality match-up scheduling problem with two and three criticality levels. *arXiv preprint arXiv:1610.07384*, 2016.
- [361] J. Nowotsch, M. Paulitsch, D. Böhler, H. Theiling, S. Wegener, and M. Schmidt. Multi-core interference-sensitive WCET analysis leveraging runtime resource capacity enforcement. In *Proc. ECRTS*, pages 109–118, 2014.
- [362] J. Nowotsch, M. Paulitsch, A. Henrichsen, W. Pongratz, and A. Schacht. Monitoring and wcet analysis in cots multi-core-soc-based mixed-criticality systems. In *Design, Automation Test in Europe Conference Exhibition (DATE)*, pages 1–5, 2014.
- [363] R. Obermaisser, Z. Owda, M. Abuteir, H. Ahmadian, and D. Weber. End-to-end real-time communication in mixed-criticality systems based on networked multicore chips. In *Proc 17th Euromicro Conference on Digital Systems Design*, pages 293–302. IEEE, 2014.
- [364] R. Obermaisser and D. Weber. Architectures for mixed-criticality systems based on networked multi-core chips. In *Proc. ETFA*, pages 1–10, 2014.
- [365] S. Osmolovskiy, I. Fedorov, V. Vinogradov, E. Ivanova, and D. Shakurov. Mixed-criticality scheduling in real-time multiprocessor systems. In *Proc. Conference of Open Innovations Association and Seminar on Information Security and Protection of Information Technology (FRUCT-ISPIT)*, pages 257–265, 2016.
- [366] A. Paolillo, P. Rodriguez, V. Svoboda, O. Desenfans, J. Goossens, B. Rodriguez, S. Girbal, M. Faugre, and P. Bonnot. Porting a safety-critical industrial application on a mixed-criticality enabled real-time operating system. In Agrawal and Easwaran, editors, *Proc. 5th Workshop on Mixed Criticality Systems (WMC)*, RTSS, pages 1–6, 2017.
- [367] A.V. Papadopoulos, E. Bini, S. Baruah, and A. Burns. AdaptMC: A Control-Theoretic Approach for Achieving Resilience in Mixed-Criticality Systems. In Sebastian Altmeyer, editor, *30th Euromicro Conference on Real-Time Systems (ECRTS 2018)*, volume 106 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 14:1–14:22. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018.
- [368] E. Papastefanakis, X. Li, and L. George. A mixed criticality approach for the security of critical flows in a network-on-chip. *ACM SIGBED Review*, 13(4):67–72, 2016.
- [369] T. Park and S Kim. Dynamic scheduling algorithm and its schedulability analysis for certifiable dual-criticality systems. In *Proc. ACM EMSOFT*, pages 253–262, 2011.
- [370] R.M. Pathan. Schedulability analysis of mixed criticality systems on multiprocessors. In *Proc. of ECRTS*, pages 309–320, 2012.
- [371] R.M. Pathan. Fault-tolerant and real-time scheduling for mixed-criticality systems. *Journal of Real-Time Systems*, 50(4):509–547, 2014.
- [372] R.M. Pathan. Improving the quality-of-service for scheduling mixed-criticality systems on multiprocessors. In Marko Bertogna, editor, *Proc. Euromicro Conference on Real-Time Systems (ECRTS)*, volume 76 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 19:1–19:22. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017.
- [373] R.M. Pathan. Improving the Schedulability and Quality of Service for Federated Scheduling of Parallel Mixed-Criticality Tasks on Multiprocessors. In Sebastian Altmeyer, editor, *30th*

- Euromicro Conference on Real-Time Systems (ECRTS 2018)*, volume 106 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 12:1–12:22. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018.
- [374] M. Paulitsch, O.M. Duarte, H. Karray, K. Mueller, D. Muench, and J. Nowotsch. Mixed-criticality embedded systems—a balance ensuring partitioning and performance. In *Proc. Euromicro Conference on Digital System Design (DSD)*, pages 453–461. IEEE, 2015.
  - [375] P. Pedro and A. Burns. Schedulability analysis for mode changes in flexible real-time systems. In *10th Euromicro Workshop on Real-Time Systems*, pages 172–179. IEEE Computer Society, 1998.
  - [376] R. Pellizzoni, P. Meredith, M-Y. Nam, M. Sun, M. Caccamo, and L. Sha. Handling mixed-criticality in soc-based real-time embedded systems. In *Proc. of the 7th ACM international conference on Embedded software, EMSOFT*, pages 235–244. ACM Press, 2009.
  - [377] R. Pellizzoni, A. Schranzhofery, J. Cheny, M. Caccamo, and L. Thiele. Worst case delay analysis for memory interference in multicore systems. In *Design, Automation Test in Europe Conference Exhibition (DATE)*, pages 741–746, 2010.
  - [378] H. Pérez, J.J. Gutiérrez, S. Peiró, and A. Crespo. Distributed architecture for developing mixed-criticality systems in multi-core platforms. *Journal of Systems and Software*, 123:145–159, 2017.
  - [379] J. Perez, D. Gonzalez, S. Trujillo, T. Trapman, and J. M. Garate. A safety concept for a wind power mixed criticality embedded system based on multicore partitioning. In *Proc. 1st WMC, RTSS*, pages 25–30, 2013.
  - [380] P. Petrakis, M. Abuteir, M.D. Grammatikakis, K. Papadimitriou, R. Obermaisser, Z. Owda, A. Papagrighoriou, M. Soulie, and M. Coppola. On-chip networks for mixed-criticality systems. In *Proc. Application-specific Systems, Architectures and Processors (ASAP)*, pages 164–169. IEEE, 2016.
  - [381] P. Pop, L. Tsiopoulos, S. Voss, O. Slotosch, C. Ficek, U. Nyman, and A. Ruiz. Methods and tools for reducing certification costs of mixed-criticality applications on multi-core platforms: the RECOMP approach. In *Proceedings of the Conference on Design, Automation and Test in Europe, WICERT, DATE*, 2013.
  - [382] T. Pop, P. Eles, and Z. Peng. Holistic scheduling and analysis of mixed time/event-triggered distributed embedded systems. In *Proceedings of the tenth international symposium on Hardware/software codesign, CODES '02*, pages 187–192. ACM, 2002.
  - [383] S. Punnekkat, R.I Davis, and A. Burns. Sensitivity analysis of real-time task sets. In *Proc. of the Conference of Advances in Computing Science - ASIAN '97*, pages 72–82. Springer, 1997.
  - [384] S. Ramanathan and A. Easwaran. MC-Fluid: rate assignment strategies. In L. Cucu-Grosjean and R. Davis, editors, *Proc. 3rd Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 6–11, 2015.
  - [385] S. Ramanathan and A. Easwaran. Utilization difference based partitioned scheduling of mixed-criticality systems. In *Design, Automation Test in Europe Conference Exhibition (DATE)*, pages 238–243, 2017.
  - [386] S. Ramanathan, A. Easwaran, and H. Cho. Multi-rate fluid scheduling of mixed-criticality systems on multiprocessors. *Real-Time Systems*, 54:247–277, 2018.
  - [387] S. Ramanathan, X. Gu, and A. Easwaran. The feasibility analysis of mixed-criticality systems. In *Proc. RTOPS, ECRTS*, 2016.

- [388] E.A. Rambo and R. Ernst. Replica-aware co-scheduling for mixed-criticality. In Marko Bertogna, editor, *Proc. Euromicro Conference on Real-Time Systems (ECRTS)*, volume 76 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 20:1–20:20. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017.
- [389] S. Ramos-Thuel and J.P. Lehoczky. Algorithms for scheduling hard aperiodic tasks in fixed priority systems using slack stealing. In *Proc. 15th IEEE Real-Time Systems Symposium*, pages 22–35, 1994.
- [390] J. Real and A. Crespo. Mode change protocols for real-time systems: A survey and a new protocol. *Journal of Real-Time Systems*, 26(2):161–197, 2004.
- [391] F. Reghenzani, G. Massari, and W. Fornaciari. Mixed time-criticality process interferences characterization on a multicore linux system. In *Proc. Euromicro Conference on Digital System Design (DSD)*, pages 427–434, 2017.
- [392] P. Regnier, G. Lima, E. Massa, G. Levin, and S. Brandt. RUN: Optimal multiprocessor real-time scheduling via reduction to uniprocessor. In *Real-Time Systems Symposium (RTSS)*, pages 104–115. IEEE, 2011.
- [393] J. Ren and L.T.X. Phan. Mixed-criticality scheduling on multiprocessors using task grouping. In *Proc. 27th ECRTS*, pages 25–36. IEEE, 2015.
- [394] P. Rodriguez, L. George, Y. Abdeddaim, and J. Goossens. Multi-criteria evaluation of partitioned EDF-VD for mixed criticality systems upon identical processors. In *Proc. 1st WMC, RTSS*, pages 49–54, 2013.
- [395] RTCA-EUROCAE. *Software Considerations in Airborne Systems and Equipment Certification DO-178B/ED-12B*. RTCA, Inc, December 1992.
- [396] S. Saewong, R. Rajkumar, J.P. Lehoczky, and M.H. Klein. Analysis of hierarchical fixed-priority scheduling. In *Proc. of the 14th Euromicro Conference on Real-Time Systems (ECRTS)*, pages 173–181, 2002.
- [397] S. Saidi, R. Ernst, S. Uhrig, H. Theiling, and B.D. de Dinechin. The shift to multicores in real-time and safety-critical systems. In *Proc. 10th International Conference on Hardware/Software Codesign and System Synthesis*, pages 220–229. IEEE Press, 2015.
- [398] M. Saksena and Y. Wang. Scaleable real-time systems design using preemption thresholds. In *Proceeding 21st IEEE Real-Time Systems Symposium.*, pages 25–34, 2000.
- [399] E. Salazar, A. Alejandro, and J. Garrido. Mixed-criticality design of a satellite software system. In *Proc. 19th World Congress The Federation of Automatic Control*, pages 12278–12283, 2014.
- [400] L. Santinelli, D. Doose, G. Durrieu, F. Boniol, C. Lesire-Cabaniols, and C. Grand. Schedulability analysis for mixed critical cyber physical systems. In *IEEE Industrial Cyber-Physical Systems (ICPS)*, pages 297–303, 2018.
- [401] L. Santinelli and L. George. Probabilities and mixed-criticalities: the probabilistic C-Space. In L. Cucu-Grosjean and R. Davis, editors, *Proc. 3rd Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 30–35, 2015.
- [402] R. Santos, S. Venkataraman, A. Das, and A. Kumar. Criticality-aware scrubbing mechanism for SRAM-based FPGAs. Technical report, Nanyang Technological University, Singapore, 2014.
- [403] J. A. Santos-Jr, G. Lima, and K. Bletsas. Considerations on the least upper bound for mixed-criticality real-time systems. In *5th Brazilian Symposium on Computing Systems Engineering (SBESC)*, 2015.



- [404] F. Santy, L. George, P. Thierry, and J. Goossens. Relaxing mixed-criticality scheduling strictness for task sets scheduled with FP. In *Proc. of the Euromicro Conference on Real-Time Systems*, pages 155–165, 2012.
- [405] F. Santy, G. Raravi, G. Nelissen, V. Nelis, P. Kumar, J. Goossens, and E. Tovar. Two protocols to reduce the criticality level of multiprocessor mixed-criticality systems. In *Proc. RTNS*, pages 183–192. ACM, 2013.
- [406] R. Schneider, D. Goswami, A. Masrur, M. Becker, and S. Chakraborty. Multi-layered scheduling of mixed-criticality cyber-physical systems. *Journal of Systems Architecture*, 59(10, Part D):1215 – 1230, 2013.
- [407] V. Sciandra, P. Courbin, and L. George. Application of mixed criticality scheduling model to intelligent transportation systems architecture. In *Proc. WiP, RTSS*, page 11, 2012.
- [408] Y. Seddik and Z. Hanzlek. Match-up scheduling of mixed-criticality jobs: Maximizing the probability of jobs execution. *European Journal of Operational Research*, 262(1):46 – 59, 2017.
- [409] L. Sha. Resilient mixed criticality systems. *CrossTalk The Journal of Defense Software Engineering*, pages 9–14, October 2009.
- [410] L. Sha, J.P. Lehoczky, and R. Rajkumar. Solutions for some practical problems in prioritizing preemptive scheduling. In *Proc. 7th IEEE Real-Time Systems Symposium*, 1986.
- [411] L. Sha, J.P. Lehoczky, and R. Rajkumar. Task scheduling in distributed real-time systems. In *Proc. of IEEE Industrial Electronics Conference*, 1987.
- [412] L. Sha, R. Rajkumar, J. Lehoczky, and K. Ramamritham. Mode change protocols for priority-driven preemptive scheduling. *Journal of Real-Time Systems*, 1(3):244–264, 1989.
- [413] L. Sha, R. Rajkumar, and J.P. Lehoczky. Priority inheritance protocols: An approach to real-time synchronisation. *IEEE Transactions on Computers*, 39(9):1175–1185, 1990.
- [414] Z. Shi and A. Burns. Real-time communication analysis for on-chip networks with wormhole switching. In *Proc. of the 2nd ACM/IEEE International Symposium on Networks-on-Chip(NoCS)*, pages 161–170, 2008.
- [415] L. Sigrist, G. Giannopoulou, N. Stoimenov, P. Huang, and L. Thiele. Mapping mixed-criticality applications on multi-core architectures. In *Proc. DATE*, pages 1–6, 2014.
- [416] D. Succi, P. Poplavko, S. Bensalem, and M. Bozga. Mixed critical earliest deadline first. Technical Report TR-2012-22, Verimag Research Report, 2012.
- [417] D. Succi, P. Poplavko, S. Bensalem, and M. Bozga. Mixed critical earliest deadline first. In *Proc. Euromicro Conference on Real-Time Systems (ECRTS)*, 2013.
- [418] D. Succi, P. Poplavko, S. Bensalem, and M. Bozga. Modeling mixed-critical systems in real-time bip. In L. George and G. Lipari, editors, *Proc. ReTiMiCS, RTCSA*, pages 29–34, 2013.
- [419] D. Succi, P. Poplavko, S. Bensalem, and M. Bozga. Time-triggered mixed critical scheduler. In *Proc. WMC, RTSS*, pages 67–72, 2013.
- [420] D. Succi, P. Poplavko, S. Bensalem, and M. Bozga. Multiprocessor scheduling of precedence-constrained mixed-critical jobs. Technical Report TR-2014-11, Verimag, Research Report, 2014.
- [421] D. Succi, P. Poplavko, S. Bensalem, and M. Bozga. Time-triggered mixed-critical scheduler on single- and multi-processor platforms. In *Proc. HPCC/CSS/ICISS*, pages 684–687, 2015.

- [422] D. Succi, P. Poplavko, S. Bensalem, and M. Bozga. Time-triggered mixed-critical scheduler on single- and multi-processor platforms. Technical Report TR-2015-8, Verimag, 2015.
- [423] B. Sprunt, J. Lehoczky, and L. Sha. Exploiting unused periodic time for aperiodic service using the extended priority exchange algorithm. In *Proc. 9th IEEE Real-Time Systems Symposium*, pages 251–258, 1988.
- [424] N. Sritharan, A.M. Kaushik, M. Hassan, and H.D. Patel. Hourglass: Predictable time-based cache coherence protocol for dual-critical multi-core systems. *CoRR*, abs/1706.07568, 2017.
- [425] W. Steiner. Synthesis of static communication schedules for mixed-criticality systems. *2012 IEEE 15th International Symposium on Object/Component/Service-Oriented Real-Time Distributed Computing Workshops*, pages 11–18, 2011.
- [426] H. Su, P. Deng, D. Zhu, and Q. Zhu. Fixed-priority dual-rate mixed-criticality systems: Schedulability analysis and performance optimization. In *Proc. Embedded and Real-Time Computing Systems and Applications (RTCSA)*, pages 59–68. IEEE, 2016.
- [427] H. Su, P. Deng, D. Zhu, and Q. Zhu. Fixed-priority dual-rate mixed-criticality systems: Schedulability analysis and performance optimization. In *IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA)*, pages 59–68, 2016.
- [428] H. Su, N. Guan, and D. Zhu. Service guarantee exploration for mixed-criticality systems. In *Proc. Embedded and Real-Time Computing Systems and Applications (RTCSA)*, pages 1–10. IEEE, 2014.
- [429] H. Su and D. Zhu. An elastic mixed-criticality task model and its scheduling algorithm. In *Proceedings of the Conference on Design, Automation and Test in Europe, DATE*, pages 147–152, 2013.
- [430] H. Su, D. Zhu, and D. Mosse. Scheduling algorithms for elastic mixed-criticality tasks in multicore systems. In *Proc. RTCSA*, 2013.
- [431] A. Taherin, M. Salehi, and A. Ejlali. Reliability-aware energy management in mixed-criticality systems. *IEEE Transactions on Sustainable Computing - online*, 2018.
- [432] D. Tamas-Selicean and P. Pop. Design optimisation of mixed criticality real-time applications on cost-constrained partitioned architectures. In *Real-Time Systems Symposium (RTSS)*, pages 24–33, 2011.
- [433] D. Tamas-Selicean and P. Pop. Optimization of time-partitions for mixed criticality real-time distributed embedded systems. In *14th IEEE International Symposium on Object/Component/Service-Oriented Real-Time Distributed Computing Workshops*, pages 2–10, 2011.
- [434] D. Tamas-Selicean and P. Pop. Task mapping and partition allocation for mixed criticality real-time systems. In *IEEE Pacific Rim Int. Sym. on Dependable Computing*, pages 282–283, 2011.
- [435] D. Tamas-Selicean and P. Pop. Design optimisation of mixed criticality real-time applications on cost-constrained partitioned architectures. *ACM Transactions on Embedded Systems*, 14(3):50:1–50:29, 2015.
- [436] J. Theis and G. Fohler. Mixed criticality scheduling in time-triggered legacy systems. In *Proc. WMC, RTSS*, pages 73–78, 2013.
- [437] J. Theis, G. Fohler, and S.K. Baruah. Schedule table generation of time-triggered mixed criticality systems. In *Proc. WMC, RTSS*, pages 79–84, 2013.

- [438] A. Thekkilakattil, R. Dobrin, and S. Punnekkat. Fault tolerant scheduling of mixed criticality real-time tasks under error bursts. In *The International Conference on Information and Communication Technologies ICICT'14*. Elsevier Procedia Computer Science, 2014.
- [439] A. Thekkilakattil, R. Dobrin, and S. Punnekkat. Mixed criticality scheduling in fault-tolerant distributed real-time systems. In *Embedded Systems (ICES), 2014 International Conference on*, pages 92–97. IEEE, 2014.
- [440] A. Thekkilakattil, A. Burns, R. Dobrin, and S. Punnekkat. Mixed criticality systems: Beyond transient faults. In L. Cucu-Grosjean and R. Davis, editors, *Proc. 3rd Workshop on Mixed Criticality Systems (WMC), RTSS*, pages 18–23, 2015.
- [441] H. Thompson. Mixed criticality systems. <http://cordis.europa.eu/fp7/ict/embedded-systems-engineering/documents/sra-mixed-criticality-systems.pdf>, EU, ICT, February 2012.
- [442] K. Tindell and A. Alonso. A very simple protocol for mode changes in priority preemptive systems. Technical report, Universidad Politecnica de Madrid, 1996.
- [443] K. Tindell, A. Burns, and A. J. Wellings. Mode changes in priority preemptive scheduled systems. In *Proc. Real Time Systems Symposium*, pages 100–109, Phoenix, Arizona, 1992.
- [444] S. Tobuschat, P. Axer, R. Ernst, and J. Diemer. IDAMC: A NoC for mixed criticality systems. In *Proc. RTCSA*, 2013.
- [445] S. Tobuschat and R. Ernst. Efficient latency guarantees for mixed-criticality networks-on-chip. In *Proc. IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 113–122, 2017.
- [446] R. Trüb, G. Giannopoulou, A. Tretter, and L. Thiele. Implementation of partitioned mixed-criticality scheduling on a multi-core platform. *ACM Trans. Embed. Comput. Syst.*, 16(5s):122:1–122:21, 2017.
- [447] S. Trujillo, A. Crespo, and A. Alonso. MultiPARTES: Multicore virtualization for mixed-criticality systems. In *Digital System Design (DSD), 2013 Euromicro Conference on*, pages 260–265, 2013.
- [448] S. Trujillo, A. Crespo, A. Alonso, and J. Perez. MultiPARTES: Multi-core partitioning and virtualization for easing the certification of mixed-criticality systems. *Microprocessors and Microsystems (online version)*, 2014.
- [449] S. Vaidhun, S.A. Arefin, Z. Guo, H. Xiong, and S.K. Das. Response time in mixed-critical pervasive systems. In *Proc. IEEE International Conference on Ubiquitous Intelligence and Computing*, 2017.
- [450] M. Vanga, H. Theiling, A. Bastoni, and B.B. Brandenburg. Supporting low-latency, low-criticality tasks in a certified mixed-criticality OS. In *Proc. RTNS*, 2017.
- [451] S. Vestal. Preemptive scheduling of multi-criticality systems with varying degrees of execution time assurance. In *Proc. of the IEEE Real-Time Systems Symposium (RTSS)*, pages 239–243, 2007.
- [452] J. Vetter, J. Fanguede, K. Chappuis, and D. Raho. VOSYSVirtualNet: Low-latency inter-world network channel for mixed-criticality systems. Technical report, ResearchGate, 2018.
- [453] E. Vitali and G. Palermo. Early stage interference checking for automatic design space exploration of mixed critical systems. In *Proc. of the 9th Workshop on Rapid Simulation and Performance Evaluation: Methods and Tools, RAPIDO*, pages 3:1–3:8. ACM, 2017.
- [454] M. Völp, M. Hähnel, and A. Lackorzynski. Has energy surpassed timeliness? scheduling energy-constrained mixed-criticality systems. In *Proc. RTAS*, pages 275–284. IEEE, 2014.

- [455] M. Völpl, A. Lackorzynski, and H. Härtig. On the expressiveness of fixed priority scheduling contexts for mixed criticality scheduling. In *Proc. WMC, RTSS*, pages 13–18, 2013.
- [456] M. Völpl, M. Roitzsch, and H. Härtig. Towards an interpretation of mixed criticality for optimistic scheduling. In *21st IEEE RTAS, Work-in-Progress*, pages 15–16, 2015.
- [457] G. von der Brüggen, K-H. Chen, W-H. Huang, and J-J. Chen. Systems with dynamic real-time guarantees in uncertain and faulty execution environments. In *Proc. Real-Time Systems Symposium (RTSS)*, pages 303–314. IEEE, 2016.
- [458] P. Wagemann, T. Distler, H. Janker, P. Raffeck, and V. Sieh. A kernel for energy-neutral real-time systems with mixed criticalities. In *Proc. RTAS*, pages 1–12. IEEE, 2016.
- [459] J. Wang and H. Wang. Work-in-progress: Scheduling of graph-based end-to-end tasks for distributed multi-criticality systems. In *Proc. IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 129–132, 2017.
- [460] Y. Wang and M. Saksena. Scheduling fixed-priority tasks with preemption threshold. In *6th Real-Time Computing Systems and Applications (RTCISA)*, pages 328–335. IEEE, 1999.
- [461] A. Wasicek, C. El-Salloum, and H. Kopetz. A system-on-a-chip platform for mixed-criticality applications. In *3th IEEE International Symposium on Object/Component/Service-Oriented Real-Time Distributed Computing (ISORC)*, pages 210–216, 2010.
- [462] R. West, Y. Li, E. Missimer, and M. Danish. A virtualized separation kernel for mixed-criticality systems. *ACM Transactions on Computer Systems (TOCS)*, 34(3):8, 2016.
- [463] Q. Xiong, F. Wu, Z. Lu, and C. Xie. Extending real-time analysis for wormhole nocs. *IEEE Transactions on Computers*, 2017.
- [464] H. Xu and A. Burns. Semi-partitioned model for dual-core mixed criticality system. In *23rd International Conference on Real-Time Networks and Systems (RTNS 2015)*, pages 257–266, 2015.
- [465] C. Yao, L. Qiao, L. Zheng, and X. Huagang. Efficient schedulability analysis for mixed-criticality systems under deadline-based scheduling. *Chinese Journal of Aeronautics*, 2014.
- [466] E. Yip, M.M.Y. Kuo, D. Broman, and P.S. Roop. Relaxing the synchronous approach for mixed-criticality systems. In *Proc. Real-Time and Embedded Technology and Application Symposium (RTAS)*, pages 89–100. IEEE, 2014.
- [467] H. Yun, G. Yao, R. Pellizzoni, M. Caccamo, and L. Sha. Memory access control in multiprocessor for real-time mixed criticality. In *Proc. of ECRTS*, pages 299–308, 2012.
- [468] A. Zabus, R.I. Davis, A. Burns, and M. González Harbour. Spare capacity distribution using exact response-time analysis. In *17th International Conference on Real-Time and Network Systems*, pages 97–106, 2009.
- [469] L. Zeng, P. Huang, and L. Thiele. Towards the design of fault-tolerant mixed-criticality systems on multicores. In *Proc. Compilers, Architectures and Synthesis for Embedded Systems*, page 6. ACM, 2016.
- [470] F. Zhang and A. Burns. Analysis of hierarchical EDF preemptive scheduling. In *Proc. of IEEE Real-Time Systems Symposium (RTSS)*, pages 423–435, 2007.
- [471] F. Zhang and A. Burns. Schedulability analysis for real-time systems with EDF scheduling. *IEEE Transaction on Computers*, 58(9):1250–1258, 2008.
- [472] N. Zhang, C. Xu, J. Li, and M. Peng. A sufficient response-time analysis for mixed criticality systems with pessimistic period. *Journal of Computational Information Systems*, 11(6):1955–1964, 2015.

- [473] X. Zhang, J. Zhan, W. Jiang, Y. Ma, and K. Jiang. Design optimization of security-sensitive mixed-criticality real-time embedded systems. In L. George and G. Lipari, editors, *Proc. ReTiMiCS, RTCSA*, pages 12–17, 2013.
- [474] Q. Zhao, Z. Al-Bayati, Z. Gu, and H. Zeng. Optimized implementation of multirate mixed-criticality synchronous reactive models. *ACM Trans. Des. Autom. Electron. Syst.*, 22(2):23:1–23:25, 2016.
- [475] Q. Zhao, Z. Gu, and H. Zeng. Integration of resource synchronization and preemption-thresholds into EDF-based mixed-criticality scheduling algorithm. In *Proc. RTCSA*, 2013.
- [476] Q. Zhao, Z. Gu, and H. Zeng. PT-AMC: Integrating preemption thresholds into mixed-criticality scheduling. In *Proc. DATE*, pages 141–146, 2013.
- [477] Q. Zhao, Z. Gu, and H. Zeng. HLC-PCP: A resource synchronization protocol for certifiable mixed criticality scheduling. *Embedded Systems Letters, IEEE*, 6(1), 2014.
- [478] Q. Zhao, Z. Gu, and H. Zeng. Resource synchronization and preemption thresholds within mixed-criticality scheduling. *ACM Transactions on Embedded Computing Systems (TECS)*, 14(4):81, 2015.
- [479] Q. Zhao, Z. Gu, H. Zeng, and N. Zheng. Schedulability analysis and stack size minimization with preemption thresholds and mixed-criticality scheduling. *Journal of Systems Architecture*, 2017.
- [480] Y. Zhao and H. Zeng. An efficient schedulability analysis for optimizing systems with adaptive mixed-criticality scheduling. *Real-Time Systems*, 53(4):467–525, 2017.
- [481] L. Ziarek and E. Blanton. The Fiji MultiVM architecture. In *Proc. JTRES*. ACM, 2015.
- [482] M. Zimmer, D. Broman, C. Shaver, and E.A. Lee. FlexPRET: A processor platform for mixed-criticality systems. In *Proc. RTAS*, pages 101–110, 2014.