Task-Set Generator for Schedulability Analysis using the TACLeBench benchmark suite

Yorick De Bock
MOdelling of Systems And Internet
Communication (MOSAIC)
University of Antwerp, Belgium
yorick.debock@uantwerpen.be

Sebastian Altmeyer
Faculty of Science
University of Amsterdam, The Netherlands
altmeyer@uva.nl

Jan Broeckhove
MOdelling of Systems And Internet
Communication (MOSAIC)
University of Antwerp, Belgium
jan.broeckhove@uantwerpen.be

Peter Hellinckx
MOdelling of Systems And Internet
Communication (MOSAIC)
University of Antwerp, Belgium
peter.hellinckx@uantwerpen.be

ABSTRACT

Current real-time embedded systems evolve towards complex systems using new state of the art technologies such as multi-core processors and virtualization techniques. Both technologies requires new real-time scheduling algorithms. For uniprocessor scheduling, utilization-based evaluation methodologies are common and well-established. For multi-core systems and virtualization, evaluating and comparing scheduling techniques using the tasks' parameters is more realistic. Evaluating these different scheduling techniques requires relevant and standardised task sets. Scheduling algorithms can be evaluated on three evaluation levels: 1) by using the mathematical model of the scheduling algorithm, 2) by simulating the scheduling algorithm and 3) by implementing the algorithm on the target platform. Generating task sets is straightforward in case of the first two levels; only the parameters of the tasks are required. Evaluating and comparing scheduling algorithms on the target platform itself, however, requires real executable tasks matching the predefined standardised task sets. Generating those executable tasks is not standardized yet.

Therefore, we developed a task-set generator that generates reproducible, standardised task sets that are suitable at all levels. Beside generating the tasks' parameters, it includes an executable generator methodology that generates executables by combining publicly available benchmarks with know execution times. This paper presents and evaluates this task-set generator. The executables approximate the wanted execution time on the hardware platform.

CCS Concepts

•**Computer systems organization** → **Real-time systems;** *Embedded software;* •**General and reference** → *Empirical studies; Measurement;*

Keywords

Real-Time Systems, Embedded Systems, Schedulability Analysis, Benchmarks, Task-Set Generator, Software Tool

1. INTRODUCTION

Current evolutions of mechatronics show an exponential growth in the number of embedded systems used for control due to the shift from mechanical control towards electronic control. This evolution creates new opportunities for even more advanced software based control. On the downside, it introduces new challenges regarding safety and reliability. Both opportunities and challenges will result in more complex software and hence higher computational requirements. These requirements can be tackled by using state of the art technologies on embedded system. Multi-core processors is a common solution to increase the computational power in a single chip. Virtualization techniques is applied to handle the complexity of the software by decomposing the software in to components which can be analysis independent from each other. For real-time systems, most applications are composed of recurring tasks with periods, deadlines and execution times. This collection of tasks is called a task set. The schedulability of task sets is very important. Uniprocessor scheduling algorithms are evaluated and compared using the maximum task set utilization which can be scheduled by the algorithm. For multiprocessor scheduling algorithms, however, the complexity of the analysis methodologies increases dramatically. This is due to the concurrent execution of tasks and the indeterminism of most multi-core architectures. Virtualization introduces a two-level hierarchical scheduling structure where the traditional analysis techniques can not be applied to. The many open issues and the practical importance of these technologies make it a topic of active research.

Parameter-based analysis uses the tasks' parameters during analysis. This results in an one-to-one mapping of task-set and scheduling algorithm; the schedulability can be analysed for a specific task set using a certain scheduling algorithm. The parameter-based analysis is more realistic for multiprocessor scheduling algorithms, and at this moment the only analysis methodology for the hierarchical schedul-

ing structure in the virtualization technology. The schedulability of a task set is the key criteria for the evaluation and comparison of scheduling algorithms. Other criteria such as the number of pre-emptions, energy consumption, scheduler overhead, cache performance etc. can also be used the evaluate and compare scheduling algorithm. The schedulability and other criteria can be evaluated at an early stage of the design process of the mechatronic system. At this stage the complexity and cost are relatively low compared to later stages of the process. However, at the later stages more evaluation criteria can be evaluated and compared with other scheduling algorithms. The performance of scheduling algorithms can be evaluated on three levels:

- **Formal proof:** at the highest evaluation level, schedulability can be formaly proven by the mathematical model of the scheduling technique.

- **Simulation-based analysis:** at the second evaluation level schedulability is validated based on the simulation model of the scheduling technique. This technique simulates a task set scheduling based on specific input parameters covering the target hardware (number of cores, ...) and the simulator settings (stepsize, simulation time, ...).

- **Implementation-based analysis:** at the last evaluation level, real-time tasks are deployed on the target platform. To schedule the tasks on the target platform, a scheduler and hence a Real-Time Operating System (RTOS) are required. The scheduler uses the scheduling algorithm to define the order of execution of the tasks.

Evaluating scheduling algorithms requires input data. In this case, a set of real-time tasks is needed to evaluate and compare different scheduling mechanisms. Depending on the evaluation level the content and format of the tasks and task set differ considerably. In the first two levels, synthetic task sets are required. This implies that the task model should only include the tasks' parameters (Worst-Case Execution Time (WCET), period and deadline) of the different tasks. It is crucial that the generated values of those parameters are not biased against any scheduling algorithm. At the third level, however, the task model is extended with executable code. The code has to match the existing task parameters. The executable tasks should be created taking into account the WCET parameter of the task. The WCET of a task is expressed in time units and will by consequence differ when deployed on different target platforms. Therefore the task model should be injected by different sets of task code when deployed on different platforms. The latter feature is missing in current evaluations of scheduling algorithms making them hard to compare or reproduce across platforms. The goal is to create executable tasks, based on the synthetic task set, for a broad set of different architectures and to examine and compare the performance of the scheduling algorithms. The execution time of the tasks will be equal to the WCET of the tasks to test the worst-case scenario for the scheduler.

In this paper we present a task-set generator tool, which not only generates synthetic task sets, but also the executable tasks for a given target platform using publicly available benchmark programs. This tool can be used to generate task sets suitable to evaluate and compare scheduling algorithms at all evaluation levels using standardized, reproducible and transparent task-set generation techniques.

The rest of the paper is organized as follows. The related work on generating synthetic task sets and executable tasks is briefly reviewed in Section 2. The task-set generator tool and its different parts are discussed in Section 3. Section 4 describes an experimental evaluation of the tool. We conclude the paper and give on overview of future research in Section 5.

2. RELATED WORK

Synthetic task-set generator tools hardly ever exist as independent tools. They often occur in literature as part of larger research projects. In [2], Bini and Buttazzo introduced the UUnifast algorithm to generate task sets for single-core processors. In 2010, Emberson et al. extended the UUnifast algorithm to multi-core systems by introducing the UUnifast-Discard algorithm and the RandFixedSum algorithm [7]. Both, UUnifast and UUnifast-Discard, are adopted as the standard task generator module in various real-time scheduling simulation tools [3–5, 11]. SimSo [4] also includes the RandFixedSum algorithm as an alternative task-set generator module.

Research on task-set generation tools that generate executable tasks is rather limited. TIMES [1], by Amnell et al., is a tool specifically designed for symbolic schedulability analysis and synthesis of executable code with predictable behaviours for real-time systems. It includes a task-set generator to integrate a set of tasks provided by the user into specific hardware architectures. Starting from a set of tasks and their runtimes it will analyse schedulability, given a specific scheduling method and hardware architecture. After successful analysis it will generate wrapping code around predefined tasks to create a compilable source code project for that specific type of hardware.

In [9] Kramer et al. present a new benchmark generator methodology for automotive applications. The automotive industry is characterized by its strong intellectual property (IP) protection. This results in a lack of realistic real-world benchmark applications. Kramer et al. propose to solve this problem by creating new benchmark applications based on code snippets origination from a well protected database containing IP protected automotive applications. The tool however was not yet available at the time of writing.

Wägemann et al. proposed GenE [13], a tool to generate benchmarks for timing analysis. It combines code patterns from real-time applications that are both representative for real-time applications and sufficiently challenging to WCET analysis tools. In addition to the source code, the generator also provides the flow facts of the benchmark. Based on this information it is straightforward to derive an accurate WCET that can be used as a reference to evaluate and compare the performance of WCET analysis techniques and technologies.

3. TASK-SET GENERATOR TOOL

The tool presented in this paper will generate task sets for the evaluation of scheduling algorithms at three levels of

abstraction in the design process. For formal and simulation based analysis, synthetic task sets are generated given a global utilization range for the task sets. A task set S consists of a set of n independent real-time tasks $\{\tau_1, \tau_2, ..., \tau_n\}$. Let τ_i indicate any given task of the task set. Each task has three parameters: the WCET C, the relative deadline D and the period T. The utilization of a task set is defined as:

$$Ut = \sum_{i=0}^{n} U_i = \sum_{i=0}^{n} \frac{C_i}{T_i} \qquad (1)$$

where C_i is the WCET of task τ_i and T_i the period of task τ_i.

For implementation-based analysis, the tasks of the synthetic task sets are extended with an executable instance. These executable tasks are created by combining benchmark programs from the TACLeBench benchmark suite [8]. The benchmark suite is a collection of benchmark programs used to evaluate timing analysis tools. The task-set generator tool calculates for each task a combination of benchmark programs. The summation of the execution time of the selected benchmark programs equals the required execution time of the task (within a pre-defined error margin). The tool distinguishes two types of user-defined input parameters: *task set specific parameters* and *program specific parameters*. The former are used to generate the synthetic task sets, the latter are used to select a set of benchmarks which fits the requirements of the user and the target platform. The selected set is used to calculate the sequence of benchmark programs for each tasks of the generated synthetic task set.

The tool is structured into three major parts which can operate independently. The first part creates the synthetic task sets, the second part selects the combination of benchmark programs for each task and the last part generates the source code and the makefile to create the executable tasks.

3.1 The Synthetic Task-Set Generator

To generate task sets to evaluate and compare scheduling algorithms, the distribution of the utilization of the tasks must be unbiased towards any scheduling algorithm [7]. The generator generates periodic tasks with implicit deadlines. These are tasks with a deadline D equals to the period T which means that the the task releases a job at every time interval T. To generate the synthetic task sets, the tool uses the *task set specific parameters* as input. The following parameters are a minimum set of parameters:

- the range of the task set utilization ($[Ut_{min}, Ut_{max}]$)
- the step value between two utilizations (Ut_{step});
- the number of task sets per utilization (k);
- the number of tasks per task set (n);
- the lower and upper bound of the period (T_l, T_u);
- the level of granularity of the period (T_Δ);
- the seed value for pseudorandom generators. (s)

In a first step, a utilization U_i for each task τ_i in the task set S is defined based on the constraint that $\sum_{i=0}^{n} U_i = U_t$

where U_t is the target utilization. The UUnifast-Discard algorithm is used to generate the task-specific utilizations. The input of this algorithm is the utilization of the task set U_t and the number of tasks in the task set n.

In the second step, the period of each task is defined and the execution time is calculated for each task. The period is randomly selected between the given lower T_l and upper bound T_u. Based on the generated period T and task utilization U, the execution time C is calculated. The mininum difference between periods of two different tasks is defined as T_Δ. To evaluate the effect of an input parameter, pseudo-random generators are used to generate these values. This implies that, for example, two identical task sets, in terms of task utilization, can be generated, but with a different lower and upper bound of the period. By keeping all other parameters fixed, all confounding effects are avoided [6].

The number of generated task sets (m) depends on the number of utilization steps and the number of task sets per utilization, $m = \left(\frac{Ut_{max} - Ut_{min}}{Ut_{step}} + 1 \right) \times k$. Another advantage of the pseudorandom generator is its ability to reproduce identical tests. An identical set of task sets can be generated by other parties in the same domain when using the same input parameters. The generated synthetic task sets are used to generated the benchmark program sequence in the second part.

3.2 Benchmark Program Sequence

After the synthetic task sets are generated, a benchmark program sequence is calculated for every task. The tool uses the TACLeBench benchmark suite as a source for the benchmark programs [8]. Within the TACLe project, the benchmark programs are formatted using the same code formatting rules, this results in a main function consisting of three function calls. These functions are present in every benchmark program:

- {*benchmark program name*}_init()
- {*benchmark program name*}_main()
- {*benchmark program name*}_return()

The first function initializes the benchmark program, the second executes the main functionality and the third function returns a variable for sanity checks. Furthermore, almost all benchmarks are platform-independent and can be compiled to and evaluated on any kind of target platform. To use the benchmark programs in the task-set generator, information regarding each benchmark must be accessible for the tool. This is realized by a complementary description file for each benchmark. This description file contains all necessary information of the benchmark (location, execution time,...). Based on the information in the description file and the selection criteria given by the user, the tool checks whether the benchmark program is suitable to be used in the benchmark program sequence. The first round of selection is based on the architecture of the target hardware; if the description file includes timing information (execution cycles) of the benchmark program on the given architecture, the benchmark program is selected. This selection results in a subset of the benchmark programs which are used to calculate the benchmark program sequence for each task. By

adding more information about the benchmark programs in their description file, a more precise selection of benchmark programs is possible. To prevent non-reproducible behaviour, each benchmark program should run at least a minimum number of times in a row; due to cache effects and other micro-architectural features, a program's execution time may vary strongly. Yet, when executed in sequence, the execution time eventually stabilizes. Hence, we also derive the minimum number of executions until a stable execution time occurs. This means that when a benchmark program is selected, the number of times a benchmark program is executed lies between its minimum number and infinity. A minimum number of consecutive executions is necessary to get a reproducible execution time. For each benchmark program, the minimum number of executions is different and does not only depends on the code of the benchmark program, but also on the architecture of the target hardware. The value of this minimum number of executions, has to be obtained in a measurement based approach by executing the benchmark program in a loop and varying the loop size. This minimum number of executions is included in the description file of the benchmark program. To calculate the program sequence for each task, we use an Integer Linear Programming (ILP) model [12] to describe the optimization problem. This model tries to match the sum of the benchmark program execution times with the target WCET of the task. The number of times a benchmark program is used, equals or is greater than zero. To build this model, the execution time of the benchmark programs must be known. These are calculated by dividing the number of execution cycles of a benchmark program by the clock frequency of the processor of the target platform. The model is also be aware of the minimum required number of executions of each benchmark program. This includes in the model as an initial cost function. If a benchmark program is selected, a cost c (minimum number of execution multiplied by the execution time of the benchmark program) is added once. The model is represented by the the following equations:

$$\textbf{Maximize} \quad \sum_{i=1}^{k}(c_i y_i + e_i l_i)$$

$$\textbf{Subject To} \quad \sum_{i=1}^{k}(c_i y_i + e_i l_i) \leq E, \quad i = 1,...,k$$

$$y_i = \begin{cases} 0 \text{ for } l_i = 0 \\ 1 \text{ for } l_i > 0 \end{cases} \quad i = 1,...,k$$
$$\text{(2)}$$

$$l_i \in \mathbb{Z}_{\leq 0}, \quad i = 1,...,k$$

$$y_i \in \{0,1\}, \quad i = 1,...,k$$

where $e_0,...,e_k$ is the set of execution times of the subset of benchmark programs. The number of times a benchmark program must be executed is represented by $l_0,...,l_k$, and E represents the targeted execution time of the task. The initial cost of each benchmark program is represented by $c_0,...,c_k$. The output of the ILP is for each benchmark program the value of l; if l is bigger than zero, the minimum number of executions for that benchmark program is added to l and the benchmark program is selected. The tool uses

the GNU Linear Programming Kit (GLPK) [10] to calculate the ILP. The output of this part is an XML file per task set, containing the selected benchmark programs and the number of executions of each program. Based on the selected benchmark programs and the number of executions, the source code is generated.

3.3 Generating Executable Tasks

The final part of the tool-chain generates the source code for each task and a *makefile* for every task set. It uses the task sets from Section 3.2 as input. For each task of a task set the initialization and main function of all selected benchmarks are called from within the code of the task. Afterwards a makefile is generated that compiles the tasks for the target platform. Each benchmark program is called within a *for-loop* statement, where l is the loop bound. After the code of the benchmark programs is appended, additional lines of code are inserted at the beginning and the end of the code, depending on the user requirements. This header and footer code can be added and/or changed in the template file. Listing 1 shows an example of generated source code. Compiling the source code files using the makefile, results in a set of executables which can be executed on the target hardware.

Listing 1: Example of generated code of a Task with minimal header and footer code

```
//header_begin
int task(void)
{
//header_end
    int i;
    for(i = 0; i <606; i++){
        bitonic_init();
        bitonic_main();
    }
    for(i = 0; i <103; i++){
        h264dec_init();
        h264dec_main();
    }
    for(i = 0; i <210; i++){
        ndes_init();
        ndes_main();
    }
    for(i = 0; i <308; i++){
        bitcount_init();
        bitcount_main();
    }

//footer_begin
}
//footer_end
```

In this section we introduced a task-set generator tool which can be used to generate tasks for the formal proof, analysis by simulation and for the analysis on the implementation level. In the next section, we evaluate the generated executable tasks by executing them on the target platform and compare the measured execution times to the targeted execution times of the tasks.

4. EXPERIMENTAL EVALUATION

In this section, we report on the experiments using the task-set generator and the results of these experiments. To compare the performance of the scheduling algorithm on different evaluation levels, the utilization of the generated task sets must be identical on each level. Consequently, the execution time of each task must be equal to the targeted WCET of the task as used in the the first two evaluation levels, or at least within an acceptable margin. The aim of our experiment is to examine the deviations of the task behaviour at the three evaluation levels: (i) the formal analysis, (ii) the simulation based analysis, and (iii) the analysis at the implementation level.

We have generated a number of synthetic task sets and executables using the task-set generator. To create the executables, we first need to derive the timing behaviour of the benchmark programs To this end, we have executed and measured 10 benchmark programs to obtain their execution times and their minimum required number of executions. After updating the description files, the task-set generator calculates the benchmark program sequence of each task.

4.1 Experiment Setup

We have conducted our experiments on a platform with an Intel(R) Xeon(R) CPU E5-2420 v2 processors at 2.20 GHz. Xen 4.5 was patched with the latest version of RT-Xen[1]. The guest domain was installed with a para-virtualized kernel. Dom0 is booted with two VCPUs, each pinned on a PCPU, and 4GB memory. The remaining ten cores where used to run the guest domain. The scheduling algorithm of the guest OS, patched by LITMUS[RT], is a global Earliest Deadline First (EDF) algorithm. In our experiments tasks are generated based on the *base_task* from the LITMUS[RT] library. For tracing the tasks in the *feather-trace* tool, included by LITMUS[RT], was used.

4.2 Execution Cycles of the Benchmarks

We selected 10 benchmark programs from the TACLeBench benchmark suite. Before these programs can be used to create executable tasks, the execution time of the benchmark must be known. To create reproducible task times, a minimum number of executions is required for each benchmark program (Section 3.2). For this, we analyse the execution time of a benchmark program by executing the benchmarks a statistically relevant number of times l. Typically for each benchmark we do $l = \{10, 50, 100, 500, 1000, 2000\}$ measurements. Besides the execution times, we calculate the mean execution time T. Based on the above measurements a value l' exists: $\forall l > l' : l \cdot T \approx T_l$ with T_l the execution time of running the benchmark l times. This value l' is defined as the minimum required executions for the benchmark program. This results in an execution time T and a minimum required executions l' for each benchmark program. We analysed the execution of each benchmark program, and observed that the execution time has a standard deviation of less than 1% if the benchmark program is executed for at least l' times. To correct for the fluctuation of a task's execution time, an extra safety margin M has to be added to the number of execution cycles of the benchmark programs. We have found that a safety margin $M = 2\%$ suffices to ensure that the

---

```
<program name="ndes" path="bench/sequential/ndes">
    <parameter name="WCET" var0="x86" var1="31518"/>
    <parameter name="min_n" var0="200"/>
</program>
```
Figure 1: Example of a benchmark program description file

Figure 2: Task Set S_5

execution time does not exceed the WCET of the task. The calculated execution times and the minimum required executions for the corresponding architecture (in this case x86) are added to the description file of the benchmark program. See Figure 1.

4.3 Creating and Executing Tasks

After updating the description files, we generated the synthetic task sets and the executable tasks. We generated 5 randomized synthetic task sets $S_1, ..., S_5$, each with 20 tasks. For each task we executed the ILP program and generated the source code based on the benchmark sequence of the task. For our experiment, we have compiled the tasks as shared libraries to call the task function (Listing 1) in each job of the real-time task in LITMUS[RT]. Because the focus of the experiment lies on measuring the execution time of tasks, we used one guest with one dedicated core and pinned the virtual cpu to a physical cpu and execute the tasks one by one. The runtime of the experiment per task is 10 seconds and we repeat the experiment 10 times. Since the aim of the experiment is deriving the execution times of the tasks, the period of each task is fixed to 1 second, this gives us an equal number of measurements for each task. This results in 100 measurements for every task of the 5 task sets.

4.4 Results

The experiments shows that the calculated WCET of the ILP program of each task has a deviation of less than 0.00001% compared the target WCET generated in the first part of the task-set generator. Comparing the execution time of the tasks on the target platform to the target WCET demonstrates that the execution time does not exceed the target WCET of the tasks. Secondly, the task with the lowest minimum execution time is task τ_9 of task set S_5, Figure 2, with an execution time of 96.4% of the target WCET. Decreasing the safety margin M would result in an higher lower bound of the execution time. The upper bound of the execution time, however, would exceed the WCET and could result in an overload situation.

5. CONCLUSION AND FUTURE WORK

This paper presents a next generation task-set generator tool. It generates reproducible task sets for the three evaluation levels to evaluate and compare scheduling algorithms for state of the art technologies. For the first two levels, synthetic task sets are generated that do not bias against any scheduling algorithm. Pseudorandom generated values make it possible to generate reproducible task sets. For the third level, a new task set generator methodology to create executable tasks for measurement-based analysis has been introduced. Using publicly available benchmark programs, making test sets reproducible and also making this tool open source enables the possibility to standardize the proposed methodology.We have evaluated the task-set generator tool in a number of experiments. They establish that the execution time of the tasks do not exceed the WCET of the task, and has a lower bound of 96.4% of the WCET. The development of the task-set generator tool is a ongoing process. At this moment a first version of this tool is publicly available under the GPL license. This version of the tool can be used to generate task sets with executable tasks for the x86 architecture using 10 benchmark programs.

As this is the first version of a task-set generator there is room for many improvements and extensions. We will focus on the two topics with the highest priority:

1. Synthetic task-set generator:

   - Extend the number of supported task models. The concept is that users should be able to tune the input parameters in a sense that the generated task sets support their research goals.

   - Extend the number of supported task set generator algorithms. The intention is to create a framework in which current but also future algorithms can be plugged in.

2. Benchmark selection:

   - Improve the method to create a stable execution time due to cache effects. This would result more stable execution times, and would decrease the safety margin M while not exceeding the targeted WCETs of the tasks.

   - Increase the number of processor architectures. This would gave us and other researchers the possibility to evaluate and compare scheduling algorithms on different target platforms using identical synthetic task sets.

   - Make the selection of the benchmark programs not only based on the architecture of the target platform, but also on other criteria. After pre-selecting the benchmark programs with support of the chosen architecture, the user will be able to select an adjusted set of benchmarks. Potential selection criteria are the use of floating point units, the size of the benchmarks, or the domain of the benchmarks.

6. ACKNOWLEDGMENT

7. REFERENCES

[1] T. Amnell, E. Fersman, L. Mokrushin, P. Pettersson, and W. Yi. TIMES: A Tool for Schedulability Analysis and Code Generation of Real-Time Systems. In *Formal Modeling and Analysis of Timed Systems (FORMATS)*, pages 60–72. Springer Berlin Heidelberg, 2003.

[2] E. Bini and G. C. Buttazzo. Measuring the performance of schedulability tests. *Real-Time Systems*, 30(1-2):129–154, 2005.

[3] Y. Chandarli, F. Fauberteau, D. Masson, S. Midonnet, and M. Qamhieh. YARTISS : A Tool to Visualize, Test, Compare and Evaluate Real-Time Scheduling Algorithms. In *Proceedings of the 3rd International Workshop on Analysis Tools and Methodologies for Embedded and Real-time Systems (WATERS)*, pages 1–12, 2012.

[4] M. Chéramy, P.-E. Hladik, and A.-M. Déplanche. SimSo : A Simulation Tool to Evaluate Real-Time Multiprocessor Scheduling Algorithms. In *Proceedings of the 5th International Workshop on Analysis Tools and Methodologies for Embedded and Real-time Systems (WATERS)*, pages 37–42, 2014.

[5] P. Courbin and L. George. FORTAS : Framework fOr Real-Time Analysis and Simulation. In *Proceedings of the 2nd International Workshop on Analysis Tools and Methodologies for Embedded and Real-Time Systems (WATERS)*, pages 21–26, 2011.

[6] R. I. Davis and A. Burns. A Survey of Hard Real-Time Scheduling for Multiprocessor Systems. *ACM Computing Surveys*, 1(4):35, 2009.

[7] P. Emberson, R. Stafford, and R. I. Davis. Techniques for the synthesis of multiprocessor tasksets. In *Proceedings of the 1st International Workshop on Analysis Tools and Methodologies for Embedded and Real-time Systems (WATERS)*, pages 6–11, 2010.

[8] H. Falk, S. Altmeyer, P. Hellinckx, B. Lisper, W. Puffitsch, C. Rochange, R. B. S. Schoeberl, P. Waegemann, and S. Wegener. TACLeBench: A Benchmark Collection to Support Worst-Case Execution Time Research. In *16th International Workshop on Worst-Case Execution Time Analysis (WCET 2016)*, page 10, 2016.

[9] S. Kramer, D. Ziegenbein, and A. Hamann. Real World Automotive Benchmarks For Free. In *Proceedings of the 6th International Workshop on Analysis Tools and Methodologies for Embedded and Real-time Systems (WATERS)*, 2015.

[10] A. Makhorin. GNU Linear Programming Kit (GLPK), 2012.

[11] A. S. Pillai and T. B. Isha. ERTSim: An embedded real-time task simulator for scheduling. In *Proceedings of the IEEE International Conference on Computational Intelligence and Computing Research (ICCIC)*, pages 1–4, 2013.

[12] J. P. Vielma. Mixed Integer Linear Programming Formulation Techniques. *Society for Industrial and Applied Mathematics (SIAM)*, 57(1):3–57, 2015.

[13] P. Wägemann, T. Distler, T. Hönig, V. Sieh, and W. Schröder-preikschat. GenE : A benchmark generator for WCET analysis. *Open Access Series in Informatics (OASIcs)*, 47:33–43, 2015.