
Supporting Multiprocessor Resource Synchronization Protocols in RTEMS

Junjie Shi¹, Jan Duy Thien Pham¹, Malte Münch¹, Jan Viktor Hafemeister¹,
Jian-Jia Chen¹, and Kuan-Hsun Chen^{1,2}

¹TU Dortmund, Department of Computer Science, Dortmund, Germany

²University of Twente, Computer Architecture and Embedded Systems, Twente,
Netherlands

Citation: [OSPERT.2022](#)

BIB_TE_X:

```
@article{DBLP:journals/corr/abs-2104-06366,  
  author      = {Junjie Shi and  
                Jan Duy Thien Pham and  
                Malte M{"u"}nch and  
                Jan Viktor Hafemeister and  
                Jian{-}Jia Chen and  
                Kuan{-}Hsun Chen},  
  title       = {Supporting Multiprocessor Resource Synchronization Protocols in {RTEMS}},  
  journal     = {CoRR},  
  volume     = {abs/2104.06366},  
  year       = {2021},  
  url        = {https://arxiv.org/abs/2104.06366},  
  eprinttype = {arXiv},  
  eprint     = {2104.06366},  
  timestamp  = {Sun, 02 Oct 2022 15:32:12 +0200},  
  biburl     = {https://dblp.org/rec/journals/corr/abs-2104-06366.bib},  
  bibsource  = {dblp computer science bibliography, https://dblp.org}  
}
```

©2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Supporting Multiprocessor Resource Synchronization Protocols in RTEMS

Junjie Shi*, Jan Duy Thien Pham*, Malte Münch*, Jan Viktor Hafemeister*, Jian-Jia Chen* and Kuan-Hsun Chen†

*Department of Computer Science, Technische Universität Dortmund, Germany

†Department of Electrical Engineering, Mathematics and Computer Science, University of Twente, the Netherlands

E-mail: {junjie.shi, jan.pham, malte.muench, jan.hafemeister, jian-jia.chen}@tu-dortmund.de, k.h.chen@utwente.nl

Abstract—When considering recurrent tasks in real-time systems, concurrent access to shared resources can cause race conditions or data corruption. Such a problem has been extensively studied since the 1990s, and numerous resource synchronization protocols have been developed for both uni-processor and multiprocessor real-time systems, with the assumption that the operating overheads are negligible. However, such overheads may also impact the performance of different protocols depending on the practical implementation, e.g., resources are accessed locally or remotely, and tasks spin or suspend themselves when the requested resources are not available. In this paper, to show the applicability of different protocols in real-world systems, we detail the implementation of several state-of-the-art multiprocessor resource synchronization protocols in RTEMS. To study the impact of the implementation overheads, we deploy these implemented protocols on a real platform with synthetic task sets. The measured results illustrate that the developed resource synchronization protocols in RTEMS are comparable to the officially supported protocol, i.e., MrsP.

I. INTRODUCTION

In multi-tasking real-time systems, the accesses to shared resources, e.g., file, memory cell, etc., are mutually exclusive, to prevent race conditions or data corruptions. A code segment that a task accesses to the shared resource(s) is called a *critical section*, which is protected by using binary semaphores or mutex locks. That is, a task must finish its execution of the critical section before another task can access the same resource. However, the mutually exclusive executions of critical sections may cause other problems, i.e., priority inversion and deadlock, which could jeopardize the predictability of the real-time system. In order to guarantee the timeliness of a real-time system, a lot of resource synchronization protocols have been developed and analyzed since 1990s for both uni-processor and multiprocessor real-time systems.

In uni-processor real-time systems, the Priority Inheritance Protocol (PIP) and the Priority Ceiling Protocol (PCP) by Sha et al. [21], as well as the Stack Resource Policy (SRP) by Baker [4] have been widely studied. Since PIP may potentially lead to a deadlock requiring additional verification to avoid [13], PCP has been relatively common and its performance has been widely accepted. Specifically, a variant of PCP has been implemented in Ada (named Ceiling locking) and in POSIX (named Priority Protect Protocol).

Because of the increasing demand of computational power of real-time systems, multiprocessor platforms have been

widely used. A lot of multiprocessor resource synchronization protocols have been proposed and extensively studied in the domain, such as the Distributed Priority Ceiling Protocol (DPCP) [20], the Multiprocessor Priority Ceiling Protocol (MPCP) [19], the Multiprocessor Stack Resource Policy (MSRP) [14], the Flexible Multiprocessor Locking Protocol (FMLP) [5], the $O(m)$ Locking Protocol (OMLP) [7], the Multiprocessor Bandwidth Inheritance (M-BWI) [12], gEDF-vpr [2], LP-EE-vpr [3], the Multiprocessor resource sharing Protocol (MrsP) [8], the Resource-Oriented Partitioned PCP (ROP-PCP) [17], the Dependency Graph Approach (DGA) for frame-based task set [11], and its extension for periodic task set (HDGA) [26].

Although the protocols above provide the timing guarantees by bounding the worst-case response time of tasks, most of them rely on the assumption that the overheads invoked by the implementation are negligible. However, rethinking of the assumption is in fact needed. Depending on their settings, e.g., local or remote execution of critical sections, multiprocessor scheduling paradigm, and the tasks' waiting semantics, the performance of different protocols is highly relevant to the implementation. For example, under a suspension-based synchronization protocol, tasks that are waiting for access to a shared resource (i.e., the resource is locked by another task) are suspended. This strategy frees the processor so that it can be used by other ready tasks, which exploits the utilization of processor, but also increases the context switch overhead due to extra en-queue and de-queue operations for each suspension. In contrast, under a spin-based synchronization protocol, the task does not give up its privilege on the processor and has to wait by spinning on the processor until it can access the requested resource and starts its critical section, which is efficient when the critical sections are short [16].

In fact, there are only a few of the protocols have been officially supported, and there are two real-time operating systems popular in the domain: the Linux Testbed for Multiprocessor Scheduling in Real-Time Systems (LITMUS^{RT}) [9], and Real-Time Executive for Multiprocessor Systems (RTEMS) [1]. LITMUS^{RT} is an experimental platform for timing analysis mainly for academic usages. Brandenburg et al. implemented DPCP, MPCP, and FMLP [6], Catellani et al. implemented MrsP [10], and Shi et al. solidate the implementation of MrsP [24]. In addition, the recently developed DGA and its

extension for periodic tasks HDGA have been implemented by Shi et al. in [22], [23]. Alternatively, RTEMS is an open-source real-time operating system which is popular for industrial applications. RTEMS has been widely used in many fields, e.g., space flight, medical, networking, etc. However, in RTEMS, only MrsP implemented by Catellani et al. in [10], is officially supported in the upstream repository.

Therefore, we believe it is beneficial to provide comprehensively support on RTEMS with resource synchronization protocols for the related researches. Afterwards, the performance of resource synchronization protocols might be clarified by system designers, and the optimizations of implementation can also be discussed. In this work, we focus on the resource synchronization protocols which are based on (semi-) partitioned scheduling, detailed as follows:

- **Partitioned Schedule:** Each task is assigned on a dedicated processor, each processor maintains its own ready queue and scheduler. Tasks are not allowed to migrate among processors, e.g., MPCP.
- **Semi-partitioned Schedule:** Unlike the pure partitioned schedule, semi-partitioned schedule allows tasks to migrate to other processors under certain conditions. For example, in DPCP and ROP-PCP, shared resources are assigned on processors, the critical sections have to be executed on the corresponding processors, where may not be the same as the original partition of a task.

Our Contribution in a nutshell: We enhance the RTEMS with the aforementioned multiprocessor resource synchronization protocols and discuss how to revise the kernel with RTEMS Symmetric Multiprocessing (SMP) support.

- To harden the open source development, we review the SMP support of RTEMS and point out the potential pitfalls during the implementation, so that the insights can be reused on any other platforms (see Section III).
- We detail the development of three multiprocessor resource synchronization protocols, i.e., MPCP, DPCP, and FMLP, and their variants in RTEMS (see Section IV).
- To study the impact of the implementation overheads, we deploy our implementations on a real platform with synthetic task sets (see Section V). The measured overheads show that our implementation overheads are comparable to the existed implementation of MrsP, in RTEMS, which illustrates the applicability of our implementations.

The patches have been released under MIT license in [25] for RTEMS 4.12. Please note that this release branch was planned to be the latest release, but significant changes warranted to bump the major number from 4 to 5. To apply our patches to RTEMS 5, a certain adaption is additionally needed.

II. SYSTEM MODEL

We consider a task set \mathbf{T} consists of n recurrent tasks to be scheduled on M symmetric and identical (homogeneous) processors. All tasks can have multiple (non-nested) critical sections, each critical section accesses one of the Z shared resources, denoted as s_z . Each task τ_i is described by a tuple $(C_i, \mu_i, T_i, D_i, q_i)$, where:

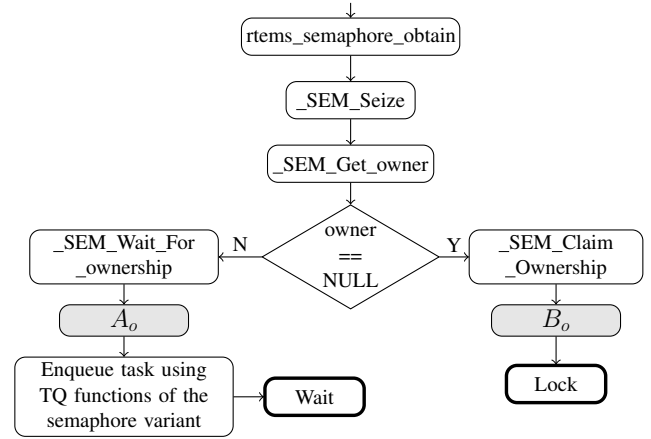


Fig. 1. Workflow of the lock directive. Block A_o and B_o are specified according to the adopted protocols.

- C_i is the worst-case execution time (WCET) of task τ_i , i.e., $C_i > 0$.
- μ_i is the set of resource(s) that τ_i requests.
- T_i is the period of task τ_i , i.e., $T_i > 0$.
- D_i is the relative deadline of the task τ_i . To fulfill its timing requirements a job of τ_i released at time t must finish its execution before its absolute deadline $t + D_i$. We consider constrained-deadline task systems, i.e., $D_i \leq T_i$ for every task $\tau_i \in \mathbf{T}$.
- q_i is the priority of task τ_i .

III. SYMMETRIC MULTIPROCESSING SUPPORT IN RTEMS

RTEMS allows users to implement new resource synchronization protocols by strictly following the RTEMS API. To create a new semaphore, `SEM_Initialize` function is called to define the specified attributes for each resource synchronization protocol. Besides the creation of semaphore, which is defined by different protocols, some common components that are similar for all the protocols, i.e., lock and unlock directives, configuration for applications, and migration mechanism, are introduced in this section.

A. Lock and Unlock Directives

The workflow of the lock directive is shown in Fig. 1. Once a task τ_i requests a shared resource, it will try to lock the corresponding semaphore. After selecting the right semaphore, denoted as `SEM`, τ_i calls the `_SEM_Seize` function. Then, the ownership of the semaphore is checked by getting the owner of the Thread queue Control. If the semaphore is locked by another task, τ_i has to wait for the owner to release the semaphore. The detailed operations in block A_o are specified according to the design of different protocols. If there is no owner yet, τ_i is set as the owner of the semaphore, and starts the execution of its critical section. The operations in block B_o can be different depending on the specified design of protocols.

The workflow of the unlock directive is shown in Fig. 2. It will be called when task τ_i has finished the execution of its critical section and releases the lock of the semaphore. The unlock directive selects the right `_SEM_Surrender` function




Fig. 2. Workflow of the unlock directive. Block A_r is specified according to the adopted protocols.

to check whether the τ_i is the current owner of the semaphore. If τ_i is not the owner, the semaphore cannot be unlocked. Otherwise, τ_i can unlock the semaphore by executing the commands in block A_r . The main function in A_r is to find the next owner for the semaphore if (at least) one task that is waiting for the semaphore. If there is no waiting task, the owner will be set to NULL accordingly. The details of the functions in A_r will be discussed in the corresponding sections for different protocols.

B. Application Configuration

In order to support semi-partitioned schedule in RTEMS, the flow for configuration in Fig. 3 has to be followed. Firstly, processors have to be bound to specific scheduler instances by using macro `_RTEMS_SCHEDULER_ASSIGN` supported in RTEMS by default. After that, each task is partitioned to a scheduler instance by using the `rtems_task_set_scheduler` directive. Each task can only be executed on the processor of the corresponding scheduler instance.




Fig. 3. The steps to configure

When a RTEMS application is configured with SMP support by following the work flow in Fig. 3, some new functions have to be implemented. In Step 1, an initial task has to be defined, which is executed in the beginning of the RTEMS application. The binding of scheduler instances to processor is based of the guide in the official c-user guide. The dedicated schedule algorithm for the scheduler instances has to be selected at first. In this paper, the *Deterministic Priority SMP Scheduler* supported in RTEMS by default is selected for all the protocols, which is the same as Fixed-Priority (FP) scheduler in the literature. Please note that, the instances have to be defined for all the available processors in the system, in order to support the semi-partitioned schedule, i.e., tasks may migrate to other processors by changing their scheduler nodes, details can be found in next subsection.

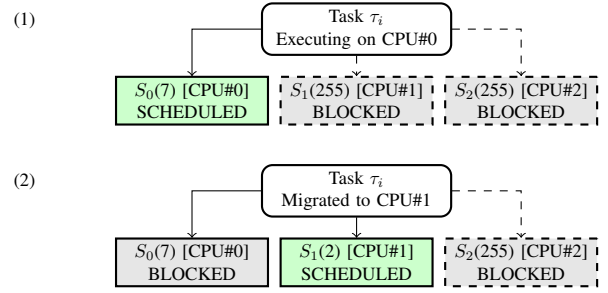


Fig. 4. Scheduler Node management: (1) Before migration, (2) After migration. Dashed blocks and lines represent that τ_i has no access to the respective scheduler instances, whereas green block is the currently used one.

C. Migration Mechanism

The migration mechanism by using arbitrary processor affinity in [15] is not supported in the current version of RTEMS. Therefore, a new migration mechanism has to be applied for those distributed-based protocols, e.g., DPCP. In our implementation, the scheduler node is modified during the run time in order to realize the task migration. When a task needs to migrate to another processor, the scheduler node of the task in its original scheduler instance is blocked, and the scheduler node of the task in its destination processor is unblocked. An additional function named `_Scheduler_Migrate_To` is implemented in `schedulerimpl.h`, which contains the task information block, the target processor, and the priority of the task in the target processor. In addition, in order to guarantee the correctness of the migration, `thread-dispatch` is disabled during the migration operation.

Fig. 4 demonstrates an example of the implemented task migration. In Fig. 4 (1), task τ_i has a scheduler node for every scheduler instance in the system. τ_i is currently executing on CPU#0 with a priority of 7 by using scheduler node S_0 , which is indicated by the the node with green background. Other two nodes with grey background are blocked, since τ_i has no access their respective scheduler instances, denoted as dashed line. In Fig. 4 (2), task τ_i performs migration to CPU#1. τ_i blocks itself on its original scheduler by using the block function of the scheduler instance on S_0 . After that, it adds S_1 to the list of its active scheduler nodes and modifies the priority of S_1 accordingly. It unblocks S_1 by using the unblock function of the corresponding scheduler instance. Migrating back to the original processor works similarly, i.e., Fig. 4 (1) is restored by using the same unblock/block function of the scheduler instances.

IV. MULTIPROCESSOR RESOURCE SYNCHRONIZATION

In this section, the implementation details of three protocols and corresponding variants are explained and discussed. Please note, we only consider non-nested resource accesses in our implementation, i.e., only one shared resource is requested during the execution of one critical section.

A. Multiprocessor Priority Ceiling Protocol

The Multiprocessor Priority Ceiling Protocol (MPCP) is a typical protocol that is based on a partitioned fixed priority (P-

Algorithm 1 MPCP implementation**Input:** Task τ_i , and `ceiling_priority` of related semaphore;

```

Function mpcp_lock():
1: if semaphore_owner is NULL then
2:   semaphore_owner  $\leftarrow$   $\tau_i$ ;
3:    $\tau_i$ .priority  $\leftarrow$  ceiling_priority;
4:    $\tau_i$  starts the execution of its critical section;
5: else
6:   Add  $\tau_i$  to the corresponding wait_queue;
7: end if

Function mpcp_unlock():
8:  $\tau_i$  releases the semaphore lock;
9: Next task  $\tau_{next} \leftarrow$  the head of the wait_queue;
10: if  $\tau_{next}$  is NULL then
11:   semaphore_owner  $\leftarrow$  NULL;
12: else
13:   semaphore_owner  $\leftarrow$   $\tau_{next}$ ;
14:    $\tau_{next}$  starts the execution of its critical section;
15: end if

```

FP) scheduler. That is, each task has a pre-defined priority, and the execution of a task is bound on a pre-defined processor, i.e., no migration is allowed. The main features of MPCP are: 1) a task will suspend itself if the resource is not available. 2) if a task is granted to access a shared resource, the priority of the task will be boosted to the ceiling priority, which equals to the highest priority of these tasks that request that resource.

The self-suspension feature is supported in RTEMS by default. In order to implement the ceiling priority boosting, one new semaphore structure is created. Besides these normal components, e.g., semaphore lock, wait queue, and current semaphore owner, one variable named `ceiling_priority` is added. Please note that, in our implementation the ceiling priority is defined by users instead of being calculated by the system dynamically. The pseudo code provided in Algo. 1 shows two main functions in our implementation, which fits the lock and unlock directive in Section III-A. The details are as follows: Once a task τ_i requests a shared resource, the ownership of the shared resource (semaphore) will be checked. If the owner of the requested shared resource is NULL, τ_i becomes the owner, and the priority of τ_i is boosted to the ceiling priority on the corresponding scheduler instance (operations in block B_o in Fig. 1). Otherwise, τ_i will be added into a wait queue, which is sorted by tasks' original priorities, i.e., task with higher priority will get earlier position (operations in block A_o in Fig. 1). Once the task τ_i finishes the execution of critical section, it will release the semaphore. The first task of the wait queue is checked, i.e., the task with the highest priority in the wait queue. If there is no task in the wait queue, the semaphore owner will be set to NULL. Otherwise, the first task of the wait queue will be set as the semaphore owner (operations in block A_r in Fig. 2).

B. Distributed Priority Ceiling Protocol

The Distributed Priority Ceiling Protocol (DPCP) is based on semi-partitioned fixed priority schedule. In DPCP, tasks

and shared resources are assigned on different processors separately, i.e., these processors that are assigned for the execution of non-critical sections are called application processors, and processors for the execution of critical sections are called synchronization processors. Once a task τ_i tries to access a shared resource, it will migrate to the corresponding synchronization processor where the shared resource is assigned on, before trying to lock the corresponding semaphore. Afterwards, these tasks on the same synchronization processor operate follow the uni-processor PCP, which been supported in RTEMS by default, i.e., Immediate Ceiling Priority Protocol (ICPP). When a task τ_i finished its execution of critical section, it will migrate back to the original application processor to continue the execution of its non-critical section, if it exists.

Hence, the main challenge of the implementation of DPCP is to allow task migrations among processors. In RTEMS, task partitioning is realized by the scheduler node in the scheduler function, i.e., scheduler node defines the original partition for each task before the execution, and stays the same during the run time. Details have been explained in Section III-C.

C. Flexible Multiprocessor Locking Protocol

In Flexible Multiprocessor Locking Protocol (FMLP), requests of shared resources are divided into two groups, i.e., long and short, according to the length of the execution time of corresponding critical section. When the requested resource is not available, a task will suspend itself if it is a long request, and a task will spin on the correspond processor if it is a short request. However, there is no conclusion regarding to how to divide requests to obtain a better schedulability. Therefore, we divided our implementation into FMLP-L which only supports long requests, and FMLP-S which only supports short requests. Please note, to simplify the implementation, all the tasks in one task set all belong to either long group or short group, no mixed division of these two groups is allowed.

In both FMLP-L and FMLP-S, the wait queue in the semaphore structure is in a FIFO order, rather than sorting by priorities like MPCP and DPCP. The operations in block B_o in Fig. 1 are as follows: In FMLP-L, we maintain a ceiling priority dynamically for each resource, which equals to the highest priority of these tasks that are currently waiting for the resource, i.e., tasks in the corresponding wait queue. The priority of the semaphore owner will be boosted to the ceiling priority if the original priority is lower than the ceiling priority, when it starts the execution of its critical section. In FMLP-S, the owner of the semaphore gets priority boosted to the highest possible priority in the system, so that the execution of its critical section is the non-preemptive. The operations in block A_o in Fig. 1 are the same for both FMLP-L and FMLP-S, i.e., add task τ_i in the end of the corresponding wait queue. The unlock operations in block A_r in Fig. 2 are also the same, i.e., try to find the next owner for the semaphore by checking the first task in the wait queue, if it exists.

Additionally, we implemented a distributed version of FMLP, denoted as DFLP, where all the requests are treated as long requests. The main difference between FMLP and DFLP

TABLE I
PROCESSOR ALLOCATION OF THE TEST APPLICATION.

CPU#0 Application	CPU#1 Application	CPU#2 Application	CPU#3 Synchronization
L (s_1)	L (s_2)	L (s_3)	-
ML (s_2)	ML (s_3)	ML (s_1)	-
M (s_3)	M (s_1)	M (s_2)	-
MH (s_2)	MH (s_3)	MH (s_1)	-
H (s_3)	H (s_1)	H (s_2)	-

is when a task requests a shared resource, it will migrate to the corresponding synchronization processor, which is similar to DPCP. The mechanism how we implement the migration has been explained in Section III-C. After the migration, critical sections are executed by following the FMLP-L on the corresponding synchronization processor(s).

V. EVALUATION AND DISCUSSION

In this section, we introduce the setup of experiments for overheads evaluation at first. Afterwards, the measured overheads are reported and analyzed. At the end, we discuss the need of formal verification over the implementation generally.

A. Experimental Setup

We evaluated the overheads of our implementations on the following platform: a NXP QorIQ T4240 RDB reference design board, which is the same as used in [10]. It has 6 GB DDR3 memory with 1866 MT/s data rate, 128 MB NOR flash(16-bit), and 2 GB SLC NAND flash. The processor T4240 contains 24-virtual-core (12 physical cores) with the PowerPC Architecture, and is running on 1.67 GHz.

To measure the overheads of our implemented protocols, timestamps are added before and after the function of our implementations. The obtain and release functions of the semaphore are measured, denoted as lock and unlock respectively. We consider a multi-processor system consists of four processors, i.e., $M = 4$, including three application processors and one synchronization processor. The total number of tasks $n = 15$, and the number of available shared resources $Z = 3$, i.e., $\mu_i \in \{s_1, s_2, s_3\}$. On each application processor, there are five tasks with five different priority levels, i.e., $q_i \in \{\text{High (H), Medium-High (MH), Medium (M), Medium-Low (ML) and Low (L)}\}$. Each task requests one of these three shared resources. Details can be found in Table I.

B. Overheads Evaluation

The overheads for different protocols are reported in Figure 5, based on more than 9,000 instances of lock and unlock operations. These distributed-based protocols, i.e., DPCP and DFLP have higher overheads than others, due to the task migrations. DFLP has the highest average overheads, since it also maintains the dynamic ceiling priority update. MrsP also has relative high overheads, since it has the help mechanism requiring task migration (however, help mechanism may not be activated all the time). Our results related to MrsP are similar

as reported in [10], i.e., 5376 ns for lock and 5514 ns for unlock on average. FMLP-L has the lowest overheads, due to the simplest mechanism. Overall, the overheads for all the protocols are relatively low and acceptable. For distributed-based protocols, we can observe that there are quite a few outliers. In fact, a similar observation has been reported in [24]. One reason could be that the behavior of cache memories kicks in to introduced operation overheads, but we have no sufficient data to pinpoint the exact cause here.

The migration overheads are measured separately, and reported in the left side of Figure 5. The results show that the overheads of task migration are significant, which might substantially affect those distributed-base protocols, i.e., DPCP and DFLP. Interestingly, we also notice that the overhead of a task to migrate to the synchronization processor is faster than migrating back to the application processor. The reason is that, normally there are more tasks running on the application processors than synchronization processors, which causes a task has to wait for longer time to obtain the scheduler instance lock on average. That is why the unlock overheads of DPCP and DFLP are higher than the lock overheads.

Although our evaluated overheads on RTEMS are similar to these protocols that are implemented on LITMUS^{RT} [10], [26], implementations of protocols on RTEMS and LITMUS^{RT} are not directly comparable due to the difference of purposes and architectures in two operating systems. Please note that RTEMS is a self-contained RTOS for real-world applications, whilst LITMUS^{RT} is a Linux-based testbed, which is mainly used for functional validation. It might be interesting to investigate which protocol is preferable on which operating systems, but it is considered out of scope here.

C. Validation and Formal Verification

To validate the correctness of our implementation, at first we test over the official coverage tests provided by RTEMS, i.e., the SMP test suites (<https://github.com/RTEMS/rtems/tree/master/testsuites/smptests>) especially, on the PowerPC device and also the QEMU emulator for ARM RealView Platform `realview-pbx-a9`, and conclude that the SMP related peripherals in RTEMS are not affected at all. Moreover, we further design several dedicated corner cases for each protocol and ensure that the designated tasks execute as the expected behaviors, which are treated as the additional coverage test for the future integration.

We note that such case-based validation may not be sufficient, since it is not possible to test over every case exhaustively. One possible way is to adopt software model checkers as proposed in [13] to detect potential data races and deadlocks in the implementation of PIP with nested locks in RTEMS. However, such searching approaches may not scale well for multiprocessor protocols unless an effective pruning strategy can be found beforehand. How to validate or formally verify an existing implementation of synchronization protocols is still an unsolved problem but out of the scope.




Fig. 5. Overheads of protocols in RTEMS (lock operation is ended by `_lk` and unlock operation is ended by `_ulkl`). The measurement of migrating a task to the synchronization processor (denoted as `mig_to`) and back to the application processor (denoted as `mig_bk`).

VI. CONCLUSION

Over the decades, quite a few number of resource synchronization protocols have been extensively studied for uniprocessor and especially multiprocessor real-time systems. In this work, we reviewed the SMP support in one popular real-time operating system RTEMS and detailed how we develop three state-of-the-art multiprocessor resource synchronization protocols, i.e., MPCP, DPCP, and FMLP, and their variants. With extensive synthetic experiments, the measured results showed that our implementations are comparable to MrsP, which is officially supported in RTEMS. Considering the real system overhead, the performance of resource synchronization protocols might be clarified and decidable by system designers.

Although several dedicated tests are provided to verify the correctness of the implementation, formal model checking is still desirable to prevent the system from potential deadlock, data races, and priority inversions. In the future work, we plan to explore on nested resource synchronization and support the arbitrary processor affinity in RTEMS to improve the generality and the efficiency. An ongoing effort is also provided to support for the latest version of RTEMS [18].

ACKNOWLEDGEMENT

This paper is supported by DFG, as part of the Collaborative Research Center SFB876, subproject A1 and A3 (<http://sfb876.tu-dortmund.de/>).

REFERENCES

- [1] RTEMS. <http://www.rtems.org/>.
- [2] B. Andersson and A. Easwaran. Provably good multiprocessor scheduling with resource sharing. *Real-Time Systems*, 46(2):153–159, 2010.
- [3] B. Andersson and G. Raravi. Real-time scheduling with resource sharing on heterogeneous multiprocessors. *Real-Time Systems*, 50(2):270–314, 2014.
- [4] T. P. Baker. Stack-based scheduling of realtime processes. *Real-Time Systems*, 3(1):67–99, 1991.
- [5] A. Block, H. Leontyev, B. Brandenburg, and J. Anderson. A flexible real-time locking protocol for multiprocessors. In *RTCSA*, pages 47–56, 2007.
- [6] B. B. Brandenburg and J. H. Anderson. An implementation of the pcp, srp, d-pcp, m-pcp, and FMLP real-time synchronization protocols in litmus^{RT}. In *RTCSA*, pages 185–194, 2008.
- [7] B. B. Brandenburg and J. H. Anderson. Optimality results for multiprocessor real-time locking. In *RTSS*, pages 49–60, 2010.
- [8] A. Burns and A. J. Wellings. A schedulability compatible multiprocessor resource sharing protocol - MrsP. In *ECRTS*, pages 282–291, 2013.
- [9] J. M. Calandrino, H. Leontyev, A. Block, U. C. Devi, and J. H. Anderson. LITMUS^{RT}: A testbed for empirically comparing real-time multiprocessor schedulers. In *RTSS*, pages 111–126, 2006.
- [10] S. Catellani, L. Bonato, S. Huber, and E. Mezzetti. Challenges in the implementation of mrsp. In *Ada-Europe*, pages 179–195, 2015.
- [11] J.-J. Chen, G. von der Brüggen, J. Shi, and N. Ueter. Dependency graph approach for multiprocessor real-time synchronization. In *IEEE Real-Time Systems Symposium, RTSS*, pages 434–446, 2018.
- [12] D. Faggioli, G. Lipari, and T. Cucinotta. The multiprocessor bandwidth inheritance protocol. In *Euromicro Conference on Real-Time Systems (ECRTS)*, pages 90–99, 2010.
- [13] S. Gadia, C. Artho, and G. Bloom. Verifying nested lock priority inheritance in RTEMS with java pathfinder. In K. Ogata, M. Lawford, and S. Liu, editors, *ICFEM*, pages 417–432, 2016.
- [14] P. Gai, G. Lipari, and M. D. Natale. Minimizing memory utilization of real-time task sets in single and multi-processor systems-on-a-chip. In *Real-Time Systems Symposium (RTSS)*, pages 73–83, 2001.
- [15] A. Gujarati, F. Cerqueira, and B. B. Brandenburg. Multiprocessor real-time scheduling with arbitrary processor affinities: from practice to theory. *Real-Time Systems*, 51(4):440–483, 2015.
- [16] G. Han, H. Zeng, M. Natale, X. Liu, and W. Dou. Experimental evaluation and selection of data consistency mechanisms for hard real-time applications on multicore platforms. *IEEE Transactions on Industrial Informatics*, 10(2):903–918, 2014.
- [17] W.-H. Huang, M. Yang, and J.-J. Chen. Resource-oriented partitioned scheduling in multiprocessor systems: How to partition and how to share? In *Real-Time Systems Symposium (RTSS)*, pages 111–122, 2016.
- [18] C. Lin, J. Shi, N. Ueter, M. Günzel, J. Reineke, and J. Chen. Type-aware federated scheduling for typed DAG tasks on heterogeneous multicore platforms. *IEEE Trans. Computers*, 72(5):1286–1300, 2023.
- [19] R. Rajkumar. Real-time synchronization protocols for shared memory multiprocessors. In *Proceedings, 10th International Conference on Distributed Computing Systems*, pages 116 – 123, 1990.
- [20] R. Rajkumar, L. Sha, and J. P. Lehoczky. Real-time synchronization protocols for multiprocessors. In *Proceedings of the 9th IEEE Real-Time Systems Symposium (RTSS '88)*, pages 259–269, 1988.
- [21] L. Sha, R. Rajkumar, and J. P. Lehoczky. Priority inheritance protocols: An approach to real-time synchronization. *IEEE Trans. Computers*, 39(9):1175–1185, 1990.
- [22] J. Shi. DGA-LITMUS-RT. <https://github.com/Strange369/Dependency-Graph-Approaches-for-LITMUS-RT>, 2018.
- [23] J. Shi. HDGA-LITMUS-RT. <https://github.com/Strange369/Dependency-Graph-Approach-for-Periodic-Tasks>, 2019.
- [24] J. Shi, K.-H. Chen, S. Zhao, W.-H. Huang, J.-J. Chen, and A. Wellings. Implementation and evaluation of multiprocessor resource synchronization protocol (mrsp) on litmus^{RT}. In *OSPERT*, 2017.
- [25] J. Shi, J. D. T. Pham, K.-H. Chen, M. Münch, J. V. Hafemeister, and J.-J. Chen. Supporting Multiprocessor Resource Synchronization Protocols in RTEMS. <https://github.com/Strange369/RTEMS-Resource-Synchronization-Protocols>, 2020.

- [26] J. Shi, N. Ueter, G. von der Brüggen, and J.-J. Chen. Multiprocessor synchronization of periodic real-time tasks using dependency graphs. In *RTAS*, pages 279–292, 2019.