

Lightweight Real-Time Synchronization under P-EDF on Symmetric and Asymmetric Multiprocessors



Alessandro Biondi^{*†} Björn B. Brandenburg^{*}

^{*}Max Planck Institute for Software Systems (MPI-SWS)

[†]Scuola Superiore Sant'Anna, Pisa, Italy

Abstract—This paper revisits lightweight synchronization under *partitioned earliest-deadline first (P-EDF) scheduling*. Four different lightweight synchronization mechanisms—namely preemptive and non-preemptive lock-free synchronization, as well as preemptive and non-preemptive FIFO spin locks—are studied by developing a new inflation-free schedulability test, jointly with matching bounds on worst-case synchronization delays. The synchronization approaches are compared in terms of schedulability in a large-scale empirical study considering both symmetric and asymmetric multiprocessors. While non-preemptive FIFO spin locks were found to generally perform best, lock-free synchronization was observed to offer significant advantages on asymmetric platforms.

I. INTRODUCTION

From a pragmatic point of view, *partitioned earliest-deadline first (P-EDF) scheduling* is in many ways an excellent choice for multiprocessor real-time systems, as it offers a favorable combination of low runtime overheads [1], good scalability [2], accurate schedulability analysis [3–5], consistently good empirical performance at high utilizations [1, 6, 7], and growing availability in commercially supported RTOS platforms (e.g., ERIKA Enterprise [8] and SCHED_DEADLINE [9] in Linux).

However, to support real-world applications, pure scheduling is not enough: applications that *share resources* such as data structures or I/O devices—which is to say, virtually any practical software system—also require *predictable* and *efficient* synchronization mechanisms.

Naturally, the question of how to best synchronize under P-EDF has received considerable attention in prior work [1, 7, 10], and a clear picture has emerged: *non-preemptive FIFO spin locks* are highly predictable, lightweight in terms of overhead, easy to use and implement, and analytically well understood [10–12]. However, as the state of the art advances, and as the hardware landscape continues to evolve, these earlier analyses and studies increasingly fail to account for important developments.

For one, none of the published studies to date has considered *asymmetric multiprocessors*. Unfortunately, as novel heterogeneous platforms—consisting of multiple processors with specialized capabilities running at different speeds (e.g., ARM’s big.LITTLE architectures or Freescale’s Vybrid platforms)—are emerging to meet the demanding design requirements of modern embedded systems, this gap in understanding is becoming progressively more noticeable. For instance, P-EDF by itself is a convenient choice for such platforms, since it allows the system designer to place each task on a core best matching its requirements (in terms of speed, instruction set

architecture, etc.). Once synchronization comes into play, though, new complications arise: when using locks, the blocking times imposed on tasks running on ‘fast’ processors will depend on critical section lengths of tasks running on ‘slow’ processors. Hence, critical sections on ‘slow’ processors can result in large blocking times compared to the timing requirements of tasks running on ‘fast’ processors. Whether spin locks are still preferable in such a setting is far from obvious.

Rather, *lock-free synchronization*, wherein the synchronization delay that tasks incur depends only on the *frequency* of conflicting accesses (and not on their *duration*), could be very attractive on asymmetric platforms. Alas, although lock-free synchronization was previously considered [7], no principled, up-to-date analysis exists in the published literature.

Another major development pertains to how synchronization delays are accounted for. In recent years, fundamentally new analysis techniques [12–14] have been developed to cope with blocking. While previous analyses [1, 7, 10] relied on the *inflation* of task parameters to model delays, it has been shown that this approach incurs substantial structural pessimism [12]. However, the new *inflation-free* analyses [12–14], which are based on *linear optimization* techniques, have been developed to date only for fixed-priority scheduling, which means that P-EDF, despite its many advantages, has regrettably fallen behind the state of the art in terms of real-time synchronization support.

This paper. Motivated by these developments, we revisit the problem of predictable and lightweight synchronization under P-EDF. First, to avoid structural pessimism, we present the first inflation-free schedulability analysis framework for P-EDF in the presence of synchronization delays (Sec. III).

Based on this foundation, we then analyze four lightweight synchronization mechanisms (Secs. IV and V) by applying the state-of-the-art approach based on linear optimization to (i) lock-free algorithms with preemptive commit loops, (ii) lock-free algorithms with non-preemptive commit loops, (iii) FIFO non-preemptive spin locks and (iv) FIFO preemptive spin locks.¹

And finally, we report on the first large-scale schedulability study (enabled by our new analyses) that investigates relative performance not only on symmetric multiprocessors (Sec. VII-A), but also on asymmetric multiprocessor platforms (Sec. VII-B) with a wide range of relative processor speeds (from 2x to 10x). Interestingly, while non-preemptive FIFO spin locks were found to still generally perform best in the symmetric case, lock-free synchronization mechanisms were observed to indeed

This paper has passed an Artifact Evaluation process. For additional details, please refer to <http://ecrts.org/artifactevaluation>.

¹Due to space constraints, only (i) and (iii) are actually shown here; the full version of this paper with all four analyses is available online [15].

```

1: procedure PUSH (Node* newNode)
2:   do
3:     oldpHead = pHead;
4:     newNode->next = oldpHead;
5:   while !CAS (pHead, newNode, oldpHead);

```

Figure 1. Example of a push operation for updating a shared lock-free queue.

offer significant advantages on asymmetric platforms.

II. BACKGROUND AND SYSTEM MODEL

In this section, we first briefly summarize the considered resource-management mechanisms, then present our system model, and finally review existing blocking-aware schedulability analysis for EDF-scheduled systems.

A. Lock-free Synchronization

Differently from lock-based approaches, lock-free algorithms do not have critical sections when accessing shared resources. The main idea is that each task works on a local copy of (a part of) the shared resource and then tries to perform an *atomic commit* to publicize its changes to the shared copy of the resource. The commit operation can fail if it is not possible to guarantee linearizability [16] of concurrent operations; in such a case, the task must *re-try* to commit its change. Lock-free algorithms are designed such that (at least) *one task always progresses* (i.e., succeeds to commit) in case of conflicts.

Figure 1 shows example code for pushing elements onto a lock-free shared queue. As it is typical for lock-free synchronization, the commit operation uses an atomic *Compare-And-Swap* (CAS) instruction, which is widely available on many current multiprocessor platforms. If a conflict occurs, the commit attempt is *repeated* until the CAS instruction succeeds.

We denote as the *commit loop* the set of instructions needed to complete the update of a shared resource (lines 2-5 in Figure 1). That is, we call the *entire* lock-free operation, including any retries, the commit loop. A commit loop correspondingly ends only when the intended update has been successfully committed. We refer to each individual iteration of a commit loop as an *attempt*, which may fail or succeed depending on the presence of conflicting attempts by other tasks.

In the absence of contention, exactly one attempt is executed. From a schedulability analysis perspective, we say that a task suffers *retry delay* when it must perform multiple attempts due to conflicts. *Arrival blocking* can occur (i.e., a preemption may be delayed) if commit loops are executed non-preemptively.

An example P-EDF schedule that demonstrates lock-free synchronization with preemptive commit loops is shown in Figure 2(a). In the example, task T_2 starts to execute first and enters a commit loop related to a shared resource ℓ . While T_2 is executing the commit loop, it is preempted by task T_1 , which commits an update to ℓ . Then, once T_1 finishes, T_2 resumes its execution. Because of the commit of T_1 , the pending commit of T_2 fails and the task must execute another commit attempt.

While T_2 is executing its second commit attempt, task T_3 (running on another processor) performs a commit to update ℓ ; hence T_2 's second attempt fails, too. These two executions of the commit loop without succeeding to commit result in retry

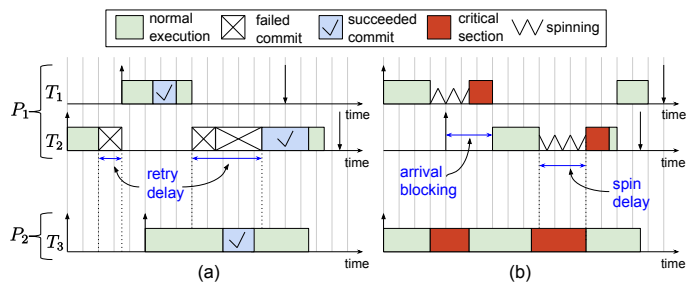


Figure 2. Example schedules of three tasks that share a resource ℓ . T_1 and T_2 execute on processor P_1 , and T_3 executes on processor P_2 . Up-arrows and down-arrows denote job releases and deadlines, respectively. (a) Example with lock-free synchronization. (b) Example with non-preemptive spin locks.

delay for T_2 . The commit loop of T_2 is finally re-executed for a third time without experiencing any conflict.

B. Spin Locks

In the case of spin locks, when a task T_i needs to access a resource that has already been locked by a task on another processor, it spins (i.e., busy-waits by wasting processor cycles) until the access to the resource is granted. Once the access is granted, the critical section using the resource is typically executed in a non-preemptive manner to ensure progress. Different types of spin locks can be identified depending on the order with which multiple requests are served, and based on whether the task remains preemptive during the spinning phase.

In this work, we consider *non-preemptive* and *preemptive FIFO spin locks*, which have been shown [12] to perform best in terms of schedulability under fixed-priority scheduling. In both types of locks, tasks wait in FIFO order. Under non-preemptive spin locks, both the spinning phase and the critical section are executed as a single non-preemptive section. Conversely, under preemptive spin locks, only the critical section is a non-preemptive section, while the task can be preempted during the spinning phase (just as during normal execution). A task that is preempted while spinning loses its position in the queue and must re-issue its lock request when it is continued.

We say that a task suffers from *spin delay* while it is busy-waiting. Spin locks can also cause arrival blocking when a needed preemption (according to EDF) is delayed due a non-preemptively executing job. Non-preemptive execution can result from (i) the execution of a critical section or, (ii) in the case of non-preemptive spin locks, because a task is busy-waiting. We denote the latter case as *transitive arrival blocking*.

An example P-EDF schedule with non-preemptive spin locks is shown in Figure 2(b). In this example, when task T_2 is released, it should be immediately scheduled according to EDF. However, it is delayed because T_1 is non-preemptively spinning while it waits for a shared resource ℓ that is currently being used by task T_3 . Once T_3 completes its critical section, T_1 accesses ℓ in a non-preemptive manner. Hence, T_2 suffers arrival blocking from its release time until the completion of the critical section of T_1 . When T_2 requests ℓ , T_3 is again using ℓ and T_2 incurs spin delay until the completion of T_3 's critical section. We refer to [12] for further details on spin locks in real-time systems.

C. System Model

We consider a set of n sporadic tasks $\tau = \{T_1, \dots, T_n\}$ scheduled on m processors P_1, \dots, P_m under P-EDF scheduling, where any ties in absolute deadline are broken in FIFO order. Each task T_i has a contention-free worst-case execution time (WCET) e_i , a minimum inter-arrival time p_i and a relative deadline $d_i \leq p_i$. A task T_i 's utilization is the fraction $u_i = e_i/p_i$; the total utilization is given by $U = \sum_{i=1}^n u_i$. Tasks do not self-suspend. We let $\tau(P_k)$ denote the set of tasks allocated to processor P_k , and let P^* denote an arbitrary processor.

The tasks share a set of n_r single-unit resources $Q = \{\ell_1, \dots, \ell_{n_r}\}$. We distinguish between two types of shared resources: **(i)** *local resources*, denoted Q^l , which are accessed *only* by tasks allocated to the same processor and **(ii)** *global resources*, denoted Q^g , which are accessed by tasks allocated on different processors, with $Q^l \cap Q^g = \emptyset$ and $Q^l \cup Q^g = Q$.

If resources are protected by locks, local resources are managed with the Stack Resource Policy (SRP) [17], while global resources are protected by FIFO spin locks [10, 12, 18] and accessed through non-preemptive critical sections. If instead resources are managed with lock-free algorithms, local and global resources are accessed using the same protocol.

For each task T_i and each resource ℓ_q , we let $N_{i,q}$ denote the maximum number of accesses (through either critical sections or commit loops), and let $L_{i,q}$ denote the respective maximum critical section (or commit loop) length. If a task T_i does not access ℓ_q we require $L_{i,q} = N_{i,q} = 0$.

If the SRP is used for managing local resources, each task T_i is assigned to a *preemption level* π_i , and each resource ℓ_q is assigned a *resource ceiling* $\pi(\ell_q)$ [17]. These parameters are computed for each processor P_k with respect to the set of tasks $\tau(P_k)$. According to the SRP, preemption levels are ordered inversely with respect to the order of relative deadlines – i.e., $\pi_i > \pi_j \Leftrightarrow d_i < d_j$ with $\{T_i, T_j\} \in \tau(P^*)$, while resource ceilings are defined as $\pi(\ell_q) = \max_{T_i \in \tau(P^*)} \{\pi_i \mid N_{i,q} > 0\}$.

We use the notation $\lceil x \rceil_0$ to denote $\max(0, \lceil x \rceil)$ and $\lfloor x \rfloor_0$ to denote $\max(0, \lfloor x \rfloor)$. Dense time is assumed.

D. EDF Analysis with Arrival Blocking

Under P-EDF scheduling, the system is schedulable if each processor is deemed schedulable by uniprocessor EDF analysis. Hence, we briefly review the analysis of EDF with arrival blocking, which is based on the notion of deadline busy-periods.

Definition 1 (deadline busy-period): An interval $[t_0, t_0 + t]$ of length t is a deadline busy-period iff **(i)** $t_0 + t$ is the absolute deadline of some job, and **(ii)** t_0 is the *last* time before $t_0 + t$ such that there are no pending jobs with a release date before t_0 and an absolute deadline before or at $t_0 + t$.

We build on Baruah's enhanced *processor demand criterion* (PDC) [19], which extends the original PDC analysis [20] to incorporate *arrival blocking* (due to either the SRP or non-preemptive sections). The original PDC analysis checks that the maximum cumulative execution demand of jobs that both are released and have an absolute deadline in any deadline busy-period of length t does not exceed the interval length t .

Consider a deadline busy-period $[t_0, t_0 + t]$. In the presence of arrival blocking, at most a *single* job with an absolute deadline past $t_0 + t$ can contribute additional processor demand [17]. The additional delay caused by such a job, which necessarily belongs to a task with a relative deadline larger than t , is accounted for in the enhanced PDC by a blocking term. According to Baruah's analysis [19], a set of n sporadic tasks is schedulable under EDF with arrival blocking if

$$\forall t \geq 0, B(t) + \sum_{i=1}^n \left\lfloor \frac{t + p_i - d_i}{p_i} \right\rfloor_0 e_i \leq t, \quad (1)$$

where $B(t)$ is an upper-bound on the maximum arrival blocking experienced in any deadline busy-period of length t .

Baruah [19] defined the blocking bound $B(t)$ as follows:

- in the case of non-preemptive blocking, $B(t) = \max\{Z_x \mid d_x > t\}$, where Z_x denotes the length of the maximum non-preemptive section of task T_x ; and
- in the case of ceiling blocking,

$$B(t) = \max\{L_{x,q} \mid d_x > t \wedge \ell_q \in pc(t)\}, \quad (2)$$

where $pc(t)$ denotes the set of local resources whose ceiling is conflicting with at least one task that has a relative deadline of length at most t , formally defined as

$$pc(t) = \{\ell_q \in Q^l \mid \pi(\ell_q) \geq \min\{\pi_x \mid d_x \leq t\}\}. \quad (3)$$

If both ceiling and non-preemptive blocking is possible, then $B(t)$ is simply the maximum of the two bounds.

From Baruah's analysis [19] we observe:

- O1** Arrival blocking in a deadline busy-period of length t is caused only by tasks with a relative deadline larger than t . In fact, under either definition, $B(t) = 0$ for $t \geq \max\{d_x\}$.
 - O2** Only a single blocking request can cause arrival blocking.
- Based on this foundation, we next present our framework for the inflation-free analysis of synchronization delays under P-EDF.

III. ANALYSIS FRAMEWORK

Our goal is to integrate **(i)** blocking times and retry delays that arise when resources are managed with lock-free algorithms, or alternatively **(ii)** blocking times that arise due to contention for resources that are protected by spin locks. For simplicity, we assume that either *all* resources are protected by locks, or that *all* resources are synchronized in a lock-free fashion.²

In contrast to the uniprocessor blocking bounds [19], we must account for both arrival blocking and the additional processor demand that results from busy-waiting or the re-execution of commit loops. However, to avoid structural pessimism [12], and in contrast to the classic MSRP analysis [10], we seek to account for this extra demand without inflating any task parameters, which introduces considerable challenges.

A. Analysis Setup Challenges

Theoretically speaking, a task set is schedulable if Eq. (1) holds for all $t \geq 0$ in an infinite and continuous domain [19].

²An analysis of hybrid setups (i.e., where there are some spin locks and some lock-free data structures) can be derived by combining our analyses of the two extremes and has been omitted due to space constraints.

To actually implement the schedulability test, two additional constraints are needed: **(i)** a discretization of the domain for the deadline busy-period length t ; and **(ii)** a maximum deadline busy-period length up to which Eq. (1) must be checked. In the presence of arrival blocking, Baruah established [19] that it is sufficient to check Eq. (1) only for deadline busy-periods of length $t \in \cup_{i=1}^n \{jp_i + d_i \mid j \in \mathbb{N}_{\geq 0}\}$ with $t \leq L$, where

$$L = \max \left\{ \max\{d_i\}, \frac{1}{1-U} \cdot \sum_{i=1}^n u_i \cdot \max\{0, p_i - d_i\} \right\}. \quad (4)$$

This bound is valid only if $U < 1$.

However, when integrating spin or retry delays, a problem arises with the maximum deadline busy-period length L . As is apparent from Eq. (4), the value of L depends on each task's utilization, which in turns depends on its WCET. When considering lock-free algorithms (resp., spin locks), the actual execution time of tasks may be greater due to retry (resp., spin) delay. Eq. (4) hence cannot be used in its original form.

Inflated [10] task utilizations, which are derived by adding a coarse-grained upper-bound on the maximum per-job delay to each task's WCET, can be computed to sidestep the problem. However, this approach can be extremely pessimistic [7, 12], and when it determines a task set to be in overload (i.e., if $U \geq 1$ after WCET inflation), Eq. (4) is invalid, which prevents the application of the PDC in its original form, even though the task set might be schedulable. To avoid such pessimism, we developed a new alternative that does not rely on inflation.

B. Inflation-Free Analysis of Synchronization Delay

The main idea is to bound the maximum deadline busy-period length using the *arrival curve* (AC) concept [21, 22]. To this end, we require the notion of a busy-period.

Definition 2 (busy-period): An interval $[t_0, t_0 + t]$ of length t is a busy-period iff, $\forall t' \in (t_0, t_0 + t)$, there is at least one job pending at time t' that was released before time t' .

The maximum busy-period length is known to upper-bound the maximum deadline busy-period length [23, 24]. In the following, we let L^+ denote a bound on the maximum busy-period length, which on processor P^* is given by the least fixed point [24, 25]:

$$L^+ = \min \left\{ t > 0 \mid \sum_{T_i \in \tau(P^*)} \left\lceil \frac{t}{p_i} \right\rceil e_i + B^{(AC)}(t) = t \right\}, \quad (5)$$

where the term $B^{(AC)}(t)$ denotes an upper-bound on the total synchronization delay (i.e., retry or spin delay) imposed on all jobs released in any busy-period of length t . Intuitively, the equation identifies the first idle time after the beginning of a busy-period starting at $t = 0$ by considering the cumulative execution time of all tasks on processor P^* .

It bears repeating that $B^{(AC)}(t)$ reflects only retry and spin delay, and not any arrival blocking: as established by Pellizzoni and Lipari [23] and Spuri [24], arrival blocking is irrelevant when bounding the maximum busy-period of sporadic tasks. (How we compute $B^{(AC)}(t)$ is discussed in Sec. VI.)

```

1: procedure PDC( $P_k, t_{LB}, t_{UB}$ )
2:   for all  $t \in \Phi_k \cap \{t_{LB} \leq t \leq t_{UB}\}$  do
3:     if  $B^{(PDC)}(t) + \sum_{T_i \in \tau(P_k)} \left\lceil \frac{t+p_i-d_i}{p_i} \right\rceil_0 e_i > t$  then
4:       return FALSE;
5:   return TRUE;

```

Figure 3. Algorithm for checking the PDC in the time window $[t_{LB}, t_{UB}]$ for the tasks allocated to processor P_k . The term $B^{(PDC)}(t)$ denotes an upper-bound on the total arrival blocking and spin (resp., retry) delay in a deadline busy-period of length t . The set Φ_k denotes the values of t at which the value of the expression in line 3 changes (i.e., the steps of the demand curve).

```

1: procedure ISCPUSCHEDULABLE( $P_k$ )
2:    $L \leftarrow \min_{T_i \in \tau(P_k)} \{e_i\}$ ;
3:   while TRUE do
4:      $L^+ \leftarrow B^{(AC)}(L) + \sum_{T_i \in \tau(P_k)} \left\lceil \frac{L}{p_i} \right\rceil e_i$ ;
5:     if  $L == L^+$  then
6:       break;
7:     if not PDC( $P_k, L, L^+$ ) then
8:       return FALSE;
9:      $L \leftarrow L^+$ ;
10:  return TRUE;

```

Figure 4. Algorithm for checking schedulability on processor P_k .

One final complication is that, if the task set is in overload, then Eq. (5) has no solution—i.e., the maximum busy-period length is unbounded. Therefore, it is not possible to directly apply Eq. (5) to bound the maximum busy-period length L^+ .

Instead, we solve Eq. (5) through an iterative fixed-point search, thereby obtaining a sequence of time intervals in which the PDC can be applied. If the processor is overloaded, then the PDC will eventually fail. Otherwise, a fixed-point iteration will converge to the maximum busy-period length.

Based on this intuition, we developed the schedulability test shown in Figures 3 and 4. The algorithm in Figure 3, which serves as the foundation for the schedulability test shown in Figure 4, checks the PDC in a given time window $[t_{LB}, t_{UB}]$.

Note that the algorithm in Figure 3 depends on the blocking term $B^{(PDC)}(t)$, which is an upper-bound on the cumulative arrival blocking and spin (resp., retry) delay in any deadline busy-period of length t , and which—in contrast to $B^{(AC)}(t)$ —does account for arrival blocking. The algorithm further depends on the set Φ_k . This set includes all points of discontinuity (i.e., the steps) of the LHS in line 3 of Figure 3. We will discuss how to compute $B^{(PDC)}(t)$ and Φ_k in Secs. IV–VI.

A system is deemed schedulable if the algorithm in Figure 4 returns TRUE for each of the processors P_1, \dots, P_m . The algorithm enters a loop in which Eq. (5) is solved through a fixed-point iteration (line 4): the loop starts from $L = \min_{T_i \in \tau(P_k)} \{e_i\}$ and iteratively produces a new *tentative* upper-bound for the busy-period length L^+ . In each iteration, the PDC is applied to the time window $[L, L^+]$, which is the interval between the last and the current tentative busy-period lengths.

The algorithm has two stop conditions: **(i)** either a potential deadline miss is identified by the PDC (line 7) and the task set is deemed *not schedulable*, or **(ii)** the fixed-point iteration converges (i.e., $L = L^+$, as tested in line 5), thus ensuring that deadline misses have been ruled out in any deadline busy-period shorter than the maximum busy-period length, which implies that all tasks in $\tau(P_k)$ are *schedulable*.

C. Blocking and Delay Computation

In the following, we present an approach for computing suitable upper-bounds $B^{(AC)}(t)$ and $B^{(PDC)}(t)$ by means of solving *optimization problems*. Specifically, we propose to obtain these bounds in a declarative fashion with either an *integer linear program* (ILP) in the case of lock-free synchronization, or a *mixed-integer linear program* (MILP) in the case of spin locks.

At a high level, our approach enumerates all possible requests for shared resources that could overlap with a given problem window and models their contribution to retry delay and/or blocking as variables of a (M)ILP, that, when maximized, yields a safe upper-bound on the maximum delay in any possible schedule. To rule out impossible scenarios, we impose constraints that encode protocol invariants (e.g., causes of retry delay, non-preemptive sections, etc.) and workload characteristics (such as task locality, request rates, etc.).

Our analysis conceptually leverages the LP-based approach first introduced by Brandenburg [13] and further developed by Wieder and Brandenburg [12]. However, these prior analyses exclusively target fixed-priority scheduling and cannot directly be applied to EDF-scheduled systems. In particular, whereas the prior analyses derive an optimization problem *for each task*, the notion of a per-task optimization problem is meaningless in our context since the PDC framework does not consider individual tasks, but rather examines deadline busy-periods *as a whole*.

Due to space constraints, we focus here on preemptable commit loops and non-preemptive spin locks; the remaining two analyses are provided online [15].

IV. DELAY DUE TO LOCK-FREE ALGORITHMS

In this section, we present our analysis of retry delay in lock-free synchronization under P-EDF, assuming that tasks remain fully preemptable at all times. To begin, let P^* be the processor under analysis and let t be an arbitrary, but fixed (i.e., constant) deadline busy-period length under observation. Moreover let $\tau^R = \tau \setminus \tau(P^*)$ denote the set of *remote tasks*.

Based on Eq. (1), we let $nljobs(T_i, t) = \lfloor (t + p_i - d_i) / p_i \rfloor_0$ denote the maximum number of jobs of a *local* task $T_i \in \tau(P^*)$ that have both release times and deadlines in an interval of length t . We further let $nrjobs(T_i, t) = \lceil (t + d_i) / (p_i) \rceil$ denote an *upper bound* on the number of jobs of a *remote* task $T_i \in \tau^R$ that are pending in any interval of length t (e.g., [1, Ch. 5]).

Since the length t is *constant*, in the context of a specific ILP instance, i.e., from the point of view of the ILP solver, $nrjobs(T_i, t)$ and $nljobs(T_i, t)$ are simply constants.

Next, we establish the key properties of lock-free algorithms that serve as the analytical basis of our ILP formulation.

A. Properties of lock-free algorithms

As explained in Sec. II-A, under lock-free synchronization, a task suffers retry delay only if it conflicts with commits of other tasks. Consider two tasks T_i and T_x that concurrently attempt to update a shared resource ℓ_q . Specifically, let c_i and c_x denote the commit loops executed by T_i and T_x , respectively.

Definition 3: We say that c_i *causes* T_x to incur retry delay if c_i has been the *last* commit to modify ℓ_q when c_x fails.

Lemma 1: A commit loop is *re-executed at most once* per each conflicting commit loop that causes retry delay.

Proof: Consider an arbitrary commit loop c_i conflicting with another arbitrary commit loop $c_x \neq c_i$. Simultaneous commits are impossible by construction in lock-free algorithms (e.g., as guaranteed by the CAS instruction). Hence, two scenarios are possible: (i) c_i commits before c_x and (ii) c_i commits after c_x . In case (i), there are no further retries of c_i since it has completed. In case (ii), c_x can cause only a *single* retry of c_i because c_x has completed. ■

Retry delay can be caused by two types of conflicts:

- *local conflicts*, which are caused by the execution of local, preempting tasks with a shorter relative deadline; and
- *remote conflicts*, which are caused by the asynchronous execution of tasks on other processors.

A key property of lock-free synchronization is that *the magnitude of retry delay does not depend on the length of conflicting commits*, but only on the *number* of local and remote conflicts. We next introduce three lemmas that characterize when and how often local and remote conflicts occur.

Lemma 2: A job incurs at most one local conflict each time that a local higher-priority job is released.

Proof: To conflict, two commit attempts must overlap in time. For local conflicts, this means that the commit attempt that suffers the conflict must have been preempted. Since jobs are sequential and since commit loops are not nested, a job executes at most one commit attempt at the time of preemption. Hence any job suffers at most one local conflict per preemption. In the absence of self-suspensions, preemptions arise only due to the release of higher-priority jobs. The claim follows. ■

Anderson et al. previously observed this property in an analysis of lock-free synchronization on uniprocessors [26].

Lemma 3: A job causes at most one local conflict per each resource that it accesses.

Proof: As in the proof of Lemma 2, observe that a commit attempt that suffers a local conflict must have been preempted. Let J be an arbitrary job and let ℓ_q be an arbitrary resource. Let \mathcal{C}_q denote the set of in-progress commit loops related to ℓ_q that have been preempted while J executes. Since tasks do not self-suspend, no commit loop in \mathcal{C}_q can make progress until J completes. Once J completes, a commit loop $c \in \mathcal{C}_q$ can be re-executed because of J ; then, c will eventually commit. Hence c will be the commit to have last updated ℓ_q and the other commit loops in $\mathcal{C}_q \setminus \{c\}$ will not retry because of J . ■

Lemma 4: A remote commit loop c_x (on any processor other than P^*) causes at most one remote conflict on processor P^* .

Proof: To conflict, two commit attempts must overlap in time. Let \mathcal{C} denote the set of commit attempts on processor P^* that overlap in time with c_x . The attempts in \mathcal{C} that succeed before the completion of c_x do not incur in a conflict with c_x . For the other attempts, an argument similar to the proof of Lemma 3 holds—once any other commit has succeeded, c_x no longer satisfies Def. 3. The claim follows. ■

In the following, we present an ILP formulation for computing the bound $B^{(PDC)}(t)$. To reduce clutter, from now on we omit

the superscript ‘(PDC)’. The modifications needed to compute the bound $B^{(AC)}(t)$ are discussed in Sec. VI-B.

B. Variables

To model retry delays in all possible schedules, we define the following variables for all tasks and resources:

- $Y_{i,j,q}^L(t) \geq 0$, an integer variable that counts the number of local conflicts incurred by T_i while accessing ℓ_q due to jobs of T_j in a deadline busy-period of length t ; and
- $Y_{i,q}^R(t) \geq 0$, an integer variable that counts the number of remote conflicts incurred by T_i while accessing ℓ_q in a deadline busy-period of length t .

If unambiguous, we omit the parameter t to improve readability.

These variables are integer because they count *events* (i.e., retries) that happened in a given schedule. The rationale behind these variables is the following.

Consider an arbitrary, but concrete schedule S , and in S consider an arbitrary deadline busy-period $[t_0, t_1]$ of length $t = t_1 - t_0$. For such a fixed trace S , it is trivial to determine $Y_{i,j,q}^L(t)$ by counting for each task T_i and each task T_j the number of times that T_i suffered local conflicts due to T_j during $[t_0, t_1]$. In other words, for *any* deadline busy-period $[t_0, t_1]$ in *any* possible schedule S , there exists a straightforward mapping from $(S, [t_0, t_1])$ to $Y_{i,j,q}^L(t)$. Further, in *any* such deadline busy-period $[t_0, t_1]$ in *any* schedule S , the total retry delay during $[t_0, t_1]$ caused by T_i ’s re-execution of commit loops that suffered local conflicts with T_j is bounded by $Y_{i,j,q}^L(t) \cdot L_{i,q}$ time units.

Since this reasoning applies to any schedule, it also applies to the (generally unknown) schedule in which the maximal retry delay is incurred. Hence, by *maximizing* the sum of the terms $Y_{i,j,q}^L(t) \cdot L_{i,q}$ for all tasks and resources, subject to constraints that express system invariants, we obtain a safe upper bound on the total worst-case delay due to local conflicts in any schedule.

Analogous reasoning applies to $Y_{i,q}^R(t)$. Based on these considerations, we arrive at the following objective.

C. Objective function

The objective of the ILP formulation is to **maximize** the delay bound $B(t)$, which is defined as

$$B(t) = \sum_{T_i \in \tau(P^*)} \sum_{\ell_q \in Q} \left(Y_{i,q}^R(t) + \sum_{T_j \in \tau(P^*)} Y_{i,j,q}^L(t) \right) \cdot L_{i,q}.$$

In the following sections, we present constraints that serve to exclude impossible schedules. The resulting attained maximum bounds the worst-case delay across all schedules *not excluded* by the constraints. Hence, a set of correct constraints yields a safe upper-bound on the maximum delay in any possible deadline busy-period of length t , including the worst-case scenario.

D. Generic constraints

The first three constraints apply to both preemptive and non-preemptive commit loops. First, a task T_i obviously incurs no retries due to a resource ℓ_q that it does not use (i.e., $N_{i,q} = 0$).

$$\textbf{Constraint 1: } \forall \ell_q \in Q, \forall T_i \in \tau(P^*) \mid N_{i,q} = 0, \\ Y_{i,q}^R + \sum_{T_j \in \tau(P^*)} Y_{i,j,q}^L = 0.$$

Similarly, a task that does not access a resource ℓ_q cannot prevent other tasks from updating ℓ_q .

$$\textbf{Constraint 2: } \forall \ell_q \in Q, \forall T_j \in \tau(P^*) \mid N_{j,q} = 0, \\ \sum_{T_i \in \tau(P^*)} Y_{i,j,q}^L = 0.$$

Next, we introduce a key constraint limiting remote conflicts.

$$\textbf{Constraint 3: } \forall \ell_q \in Q,$$

$$\sum_{T_i \in \tau(P^*)} Y_{i,q}^R \leq \sum_{T_x \in \tau^R} nrjobs(T_x, t) \cdot N_{x,q}.$$

Proof: Suppose not. Then there exists a schedule in which **(i)** at least one remote commit loop to update ℓ_q caused a local commit loop to retry more than once, or **(ii)** at least one local commit loop related to ℓ_q has been re-executed more than once due to the same remote commit loop. Case (i) is impossible by Lemma 4, while Lemma 1 rules out case (ii). Contradiction. ■

E. Constraints for preemptive commit loops

With the following constraint, we enforce that no retries for commit loops are considered from tasks that cannot have any jobs with a release time and absolute deadline both in a deadline busy-period of length t .

$$\textbf{Constraint 4: } \forall T_i \in \tau(P^*) \mid nljobs(T_i, t) = 0, \\ \sum_{\ell_q} \left(Y_{i,q}^R + \sum_{T_j \in \tau(P^*)} Y_{i,j,q}^L \right) = 0.$$

Proof: As stated in Sec. II-D, in any deadline busy-period with length t , a task T_i with $nljobs(T_i, t) = 0$ can contribute demand only because of arrival blocking. However, when using preemptive commit loops, there is no arrival blocking. ■

We now look at retries due to local conflicts.

$$\textbf{Constraint 5: } \forall T_i \in \tau(P^*), \forall T_j \in \tau(P^*),$$

$$\sum_{\ell_q} Y_{i,j,q}^L \leq \left\lceil \frac{d_i - d_j}{p_j} \right\rceil_0 \cdot nljobs(T_i, t).$$

Proof: Once T_i is released, it can only be preempted by local tasks T_j with a shorter relative deadline $d_j < d_i$.³ Suppose (w.l.o.g.) that the task T_i is released at time 0: then a task T_j will preempt T_i only if it is released *before* time $d_i - d_j$, otherwise it has a later absolute deadline than T_i . In the interval $[0, d_i - d_j]$, at most $\lceil (d_i - d_j)/p_j \rceil_0$ jobs of T_j are released. Hence, the RHS upper-bounds the number of jobs of T_j that preempt T_i in a deadline busy-period of length t . By Lemmas 2 and 3, this also bounds the maximum number of retries caused by T_j . ■

Constraint 6:

$$\forall T_j \in \tau(P^*), \forall \ell_q \in Q, \sum_{T_i \in \tau(P^*)} Y_{i,j,q}^L \leq \left\lceil \frac{t}{p_j} \right\rceil.$$

Proof: Suppose not. Then there exists a schedule in which **(i)** a commit loop related to ℓ_q performed more than one retry due to the same job of T_j or **(ii)** a job of T_j caused more than one retry of a commit loop related to ℓ_q . Lemma 2 rules out case (i), and Lemma 3 rules out case (ii). Contradiction. ■

To further bound the number of retries, we apply response-time analysis on commit loops, i.e., we bound the maximum

³Only tasks with strictly shorter relative deadlines are considered since FIFO tie-breaking is assumed.

time span during which a task executes a commit loop, lasting from the beginning of the loop until the successful completion of the commit, including any retries and times of preemption.

Definition 4: We let $W_{i,q}^P$ denote an upper bound on the *commit-loop response time*, which is defined as follows: if T_i enters a commit loop to update ℓ_q at time t_e , and if T_i completes the commit at time t_c , then $t_c - t_e \leq W_{i,q}^P$.

Due to space constraints, details on the computation of $W_{i,q}^P$ are given in the full version of the paper [15]. Once the bound $W_{i,q}^P$ has been found, we can limit the number of remote conflicts.

Constraint 7: $\forall T_i \in \tau(P^*), \forall \ell_q \in Q,$

$$Y_{i,q}^R \leq nljobs(T_i, t) \cdot N_{i,q} \cdot \sum_{T_x \in \tau^R} nrjobs(T_x, W_{i,q}^P) \cdot N_{x,q}.$$

Proof: Each update of ℓ_q performed by T_i completes in $W_{i,q}^P$ time units. Hence, for each commit loop there are at most $nrjobs(T_x, W_{i,q}^P) \cdot N_{x,q}$ requests issued by a remote task T_x that can cause remote conflicts. The constraint follows by noting that there are at most $nljobs(T_i, t) \cdot N_{i,q}$ commit loops to update ℓ_q issued by jobs of T_i with release time and absolute deadline in any deadline busy-period of length t . ■

V. BLOCKING DUE TO SPIN LOCKS

In this section, we present our blocking analysis for spin locks. While the analysis is conceptually similar to the one used in Sec. IV, spin locks have some properties that require the use of a different approach to obtain suitable bounds on blocking. In particular, differently from lock-free algorithms, the blocking imposed by spin locks depends on the critical section lengths of conflicting requests. To avoid accounting for any critical section more than once, we adopt a different modeling strategy [12, 13] that considers each critical section individually.

A. Variables

We model the impact of spin locks on blocking times in any possible schedule with the following variables for each task $T_i \in \tau$ and for each resource $\ell_q \in Q$:

- $X_{i,q}^S(t) \geq 0$, a real variable expressing the contribution of requests for ℓ_q issued by T_i to the *spin delay* incurred by jobs with release time and absolute deadline in any deadline busy-period of length t and that execute on P^* ;
- $X_{i,q}^A(t) \geq 0$, a real variable expressing the contribution of requests for ℓ_q issued by T_i to the *arrival blocking* suffered by tasks on P^* in a deadline busy-period of length t ;
- $A_q \in [0, 1]$, a binary variable such that $A_q = 1$ if and only if arrival blocking is caused by requests for ℓ_q .

The definitions of these variables differ from those defined for lock-free algorithms in Sec. IV-B, based on the following rationale: given an arbitrary, but concrete schedule with a deadline busy-period of length t , the requests of a task T_i for resource ℓ_q contribute to spin delay for $X_{i,q}^S(t) \cdot L_{i,q}$ time units and to arrival blocking for $X_{i,q}^A(t) \cdot L_{i,q}$ time units. The variables are defined as reals since they model fractions of time.

B. Objective Function

The objective is to **maximize** the blocking bound $B(t)$:

$$B(t) = \sum_{\forall T_i \in \tau} \sum_{\ell_q \in Q} (X_{i,q}^S(t) + X_{i,q}^A(t)) \cdot L_{i,q}. \quad (6)$$

The same considerations reported in Sec. IV-C also apply to the above objective function: by ruling out impossible schedules, a maximal solution of the MILP represents an upper bound on the worst-case blocking in any possible schedule. We use the same notation as in Sec. IV whenever not differently specified.

C. Generic Constraints

We begin with constraints that are applicable to every type of spin lock. Again, we begin by excluding trivial cases.

Constraint 8: $\forall T_i \in \tau(P^*) \mid d_i \leq t, \forall \ell_q \in Q, X_{i,q}^A = 0.$

Proof: Follows from Observation O1. ■

Constraint 9: $\sum_{\forall T_i \in \tau(P^*)} \sum_{\ell_q \in Q} X_{i,q}^S = 0.$

Proof: Only remote tasks cause spin delay. ■

By definition, spin delay and arrival blocking never coincide.

Constraint 10:

$$\forall T_i \in \tau^R, \forall \ell_q \in Q, X_{i,q}^S + X_{i,q}^A \leq nrjobs(T_i, t) \cdot N_{i,q}.$$

Proof: First note that the RHS expresses the maximum number of remote requests for ℓ_q in any interval of length t . Suppose that the constraint does not hold. Then there exists a schedule with a deadline busy-period of length t in which a request for ℓ_q at some point in time contributes to both arrival blocking and spin delay. However, arrival blocking occurs only due to the execution of a task $T_l \in \tau(P^*)$ with a relative deadline $> t$. In contrast, spin delay occurs due to the execution of a task $T_h \neq T_l \in \tau(P^*)$ with a relative deadline $\leq t$. T_l and T_h cannot execute simultaneously on P^* . Contradiction. ■

The next four constraints characterize arrival blocking.

Constraint 11: $\sum_{\ell_q \in Q} A_q \leq 1.$

Proof: Follows from Observation O2. ■

Recall from Sec. II-D that not all local resources can cause local blocking, depending on their resource ceilings.

Constraint 12: $\forall \ell_q \in Q^l \setminus pc(t), A_q = 0.$

Proof: Follows directly from Eq. (2). ■

Moreover, depending on the deadline busy-period length t , there may be no tasks that can cause arrival blocking due to a specific resource as captured by the following constraint.

Constraint 13: $\forall \ell_q \in Q, A_q \leq \sum_{T_i \in \tau(P^*) \mid d_i > t} N_{i,q}.$

Proof: Following Observation O1, arrival blocking can be caused only by a task T_i with a relative deadline $d_i > t$. Arrival blocking due to ℓ_q can occur only if at least one such task accesses ℓ_q (hence yielding a positive RHS). ■

Finally, we exclude scenarios where more than one local request causes arrival blocking by linking the decision variable A_q to the blocking variables for every resource ℓ_q .

Constraint 14: $\forall \ell_q \in Q, \sum_{\forall T_i \in \tau(P^*)} X_{i,q}^A \leq A_q.$

Proof: If $A_q = 0$ then (by definition of the variable A_q) it is not possible to have arrival blocking from ℓ_q . If $A_q = 1$, then the constraint follows from Observation O2. ■

D. Constraints for Non-preemptive FIFO Spin Locks (MSRP)

In this section, we present constraints specific to non-preemptive FIFO spin locks, which are the type of spin locks mandated in Gai et al.'s MSRP [10].

The most important constraint encodes the FIFO property and expresses that each local request (i.e., one issued on processor P^*) is delayed by *at most one request* per remote processor.

Constraint 15: $\forall P_k \neq P^*, \forall \ell_q \in Q,$

$$\sum_{T_x \in \tau(P_k)} X_{x,q}^S \leq \sum_{T_i \in \tau(P^*)} nljobs(T_i, t) \cdot N_{i,q}.$$

Proof: First note that the RHS of the constraint expresses the maximum number of requests issued by local tasks that can experience spin delay (i.e., those that have jobs with both a release time and absolute deadline within an interval of length t). Suppose the constraint does not hold. Then there exists a schedule in which a request issued by a task $T_i \in \tau(P^*)$ is delayed by more than one request issued by a task $T_x \in \tau(P_k)$. Since requests are served in FIFO order, and since tasks execute non-preemptively while spinning and while executing critical sections, each request can be delayed by at most one request issued by another processor. Contradiction. ■

Another bound on per-task spin delay can be derived by exploiting the periods and deadlines of conflicting tasks.

Constraint 16: $\forall T_x \in \tau^R, \forall \ell_q \in Q,$

$$X_{x,q}^S \leq nrjobs(T_x, t) \cdot \sum_{T_i \in \tau(P^*)} nrjobs(T_i, d_x) \cdot N_{i,q}.$$

Proof: A job of a remote task T_x overlaps with at most $nrjobs(T_i, d_x)$ jobs of a local task T_i . Hence, the requests for ℓ_q issued by a job of T_x conflict with at most $nrjobs(T_i, d_x) \cdot N_{i,q}$ requests of T_i . In particular, when using non-preemptive FIFO spin locks, each request for ℓ_q issued by T_i is delayed by at most one request for ℓ_q issued by T_x . Hence, every job of T_x delays jobs of a local task T_i with at most $nrjobs(T_i, d_x) \cdot N_{i,q}$ requests. The constraint follows since at most $nrjobs(T_x, t)$ jobs of T_x overlap with an interval of length t . ■

We now address arrival blocking. A remote request transitively contributes to the arrival blocking on processor P^* only if a local task T_x with a relative deadline $d_x > t$ is non-preemptively spinning while waiting for the completion of the transitively blocking request. Again, due to the FIFO ordering, at most one request per remote processor can transitively block.

Constraint 17: $\forall P_k \neq P^*, \forall \ell_q \in Q,$

$$\sum_{T_x \in \tau(P_k)} X_{x,q}^A \leq A_q.$$

Proof: If $A_q = 0$, then by definition there is no arrival blocking due to requests for ℓ_q . If $A_q = 1$, according to Observation O2 arrival blocking arises due to at most a single request r issued on P^* . Since requests are served in FIFO order, and since tasks execute non-preemptively while spinning or holding locks, at most one request issued on P_k can delay r . ■

VI. ARRIVAL CURVES AND TEST POINTS

Secs. IV and V report how to determine the PDC blocking bound $B^{(PDC)}(t)$ for use in the algorithm given in Figure 3. However, to obtain a practical schedulability test, we still require the discrete set of test points Φ_k and the arrival-curve blocking bound $B^{(AC)}(t)$ for use in the fix-point search in Figure 4.

A. PDC Evaluation Points Φ_k

To correctly implement the PDC, the set Φ_k of deadline busy-period lengths that are to be tested must include all points of discontinuity of the LHS of the equation at line 3 in Figure 3. Since this term depends on the bound $B^{(PDC)}(t)$, there are further discontinuity points in addition to those considered by Baruah [19] in the uniprocessor case.

In particular, as $B^{(PDC)}(t)$ is computed by solving (M)ILP formulations, discontinuities arise due to any dependencies of constraints on the deadline busy-period length t . The set Φ_k is hence given by the union of the following contributing sets: **(i)** $\cup_{T_i \in \tau(P_k)} \{jp_i + d_i \mid j \in \mathbb{N}_{\geq 0}\}$, which are the standard PDC check-points for sporadic tasks (as established by Baruah [19]), and which also correspond to the steps caused by the function $nljobs(T_i, t)$, the inequalities involving relative deadlines of tasks in Constraints 8 and 13, and variations in the constituency of the set $pc(t)$; **(ii)** $\cup_{T_i \in \tau \setminus \tau(P_k)} \{jp_i - d_i + \epsilon \mid j \in \mathbb{N}_{\geq 1}\}$, which reflect steps of the function $nrjobs(T_i, t)$; and **(iii)** $\cup_{T_i \in \tau(P_k)} \{jp_i + \epsilon \mid j \in \mathbb{N}_{\geq 1}\}$, which corresponds to the steps of the term $\lceil t/p_i \rceil$ in Constraint 6. In (ii) and (iii), $\epsilon > 0$ is an arbitrary, sufficiently small, positive number.

B. Computing the Bound $B^{(AC)}(t)$,

Recall from Sec. III-B that the AC bound simply considers *all* jobs released in the problem window, in contrast to the PDC, which considers *only* jobs that are both released *and* have an absolute deadline in the problem window

For this reason, a different definition of the number of relevant local jobs $nljobs(T_i, t)$ is needed: $nljobs(T_i, t) = \left\lfloor \frac{t}{p_i} \right\rfloor$. This follows directly from Eq. (5). As stated in Sec. III-B, when computing the blocking for the AC we do not have to account for arrival blocking [23]. Hence the constraint $\sum_{\ell_q \in Q} A_q \leq 0$ holds. The rest of the MILP and ILP formulations presented in the preceding sections can be reused without changes.

VII. EXPERIMENTAL EVALUATION

We implemented the proposed analyses in SchedCAT [27], which uses the *GNU Linear Programming Kit* or *CPLEX* as the underlying (M)ILP solver. To assess the new analyses, and to compare the different mechanisms in different scenarios, we carried out a large-scale schedulability study based on randomly generated synthetic workloads. For brevity, we denote FIFO non-preemptive and FIFO preemptive spin locks as F|N and F|P, respectively. As a baseline, we used the classic analysis for F|N spin locks under P-EDF scheduling by Gai et al. [10].

We conducted three experiments: two of them are discussed herein, while the third one is provided online [15].

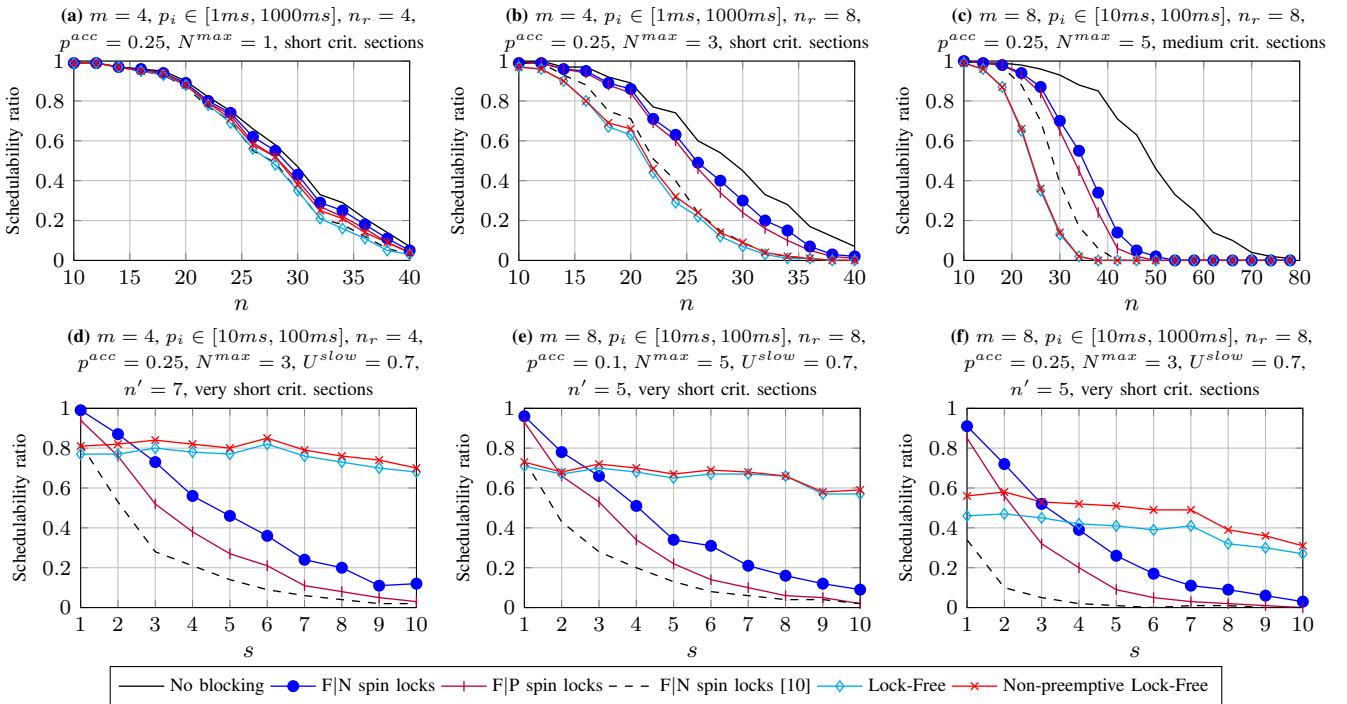


Figure 5. Experimental results of six representative configurations. Figures (a), (b) and (c) show results from Experiment 1 (Sec. VII-A) while Figures (d), (e) and (f) show results from Experiment 2 (Sec. VII-B). The values of the parameters used in each configuration are reported in the captions above the figures.

A. Experiment 1: symmetric platforms

1) *Workload generation:* We considered platforms with $m \in \{2, 4, 8\}$ identical processors. For a given n , we randomly generated task sets with the following different parameter settings. Task periods were randomly chosen from a log-uniform distribution with a range of either $[10ms, 100ms]$ or $[1ms, 1000ms]$. All tasks were assigned implicit deadlines. The utilization of each task was randomly drawn from an exponential distribution with a mean of 0.1. Each task T_i was allocated to processor P_x , where $x = \lceil i/m \rceil$. We assumed the presence of $n_r \in \{m/2, m, 2m\}$ shared resources. Each task T_i was configured to access each resource ℓ_q with a probability $p^{acc} \in \{0.1, 0.25, 0.5\}$. If T_i was chosen to access ℓ_q , then the number of requests $N_{i,q}$ was chosen uniformly at random from the set $\{1, \dots, N^{max}\}$, where N^{max} varied across $\{1, 3, 5, 7, 10\}$. The maximum critical section (or commit loop) length $L_{i,q}$ was chosen uniformly at random from either $[1\mu s, 25\mu s]$ (*short*) or $[25\mu s, 100\mu s]$ (*medium*).

For simplicity, we used the same critical section (or commit loop) length under both spin locks and lock-free alternatives. This is unlikely to be the case in practice (many lock-free approaches incur copy overhead), but it ensures that all reported trends are solely due to differences in analysis.

2) *Results:* In our study, we evaluated more than 2000 different configurations, varying the number of tasks n from 10 to $10m$ in steps of $m/2$ for each configuration. As the number of tasks increases, the overall system load and the contention for shared resources increases. For each value of n , 500 different task sets were generated and tested. Across the range of all tested configurations, we noticed that, even with the new analysis, the pessimism gap [7] between spin

locks and lock-free algorithms remains present. Once more, F|N spin locks perform best in most of the configurations and our new analysis consistently outperforms the baseline from [10], showing significant improvements especially in the presence of high contention. In some particular configurations, F|P spin locks are found to dominate all other studied mechanisms, performing slightly better than even F|N spin locks. Although specific task sets (e.g., those containing tasks with short periods) can highly benefit from the lower arrival blocking imposed under F|P spin locks, the marginal gain in schedulability observed (on average) in this study does not justify the additional complexity needed for their support [28, 29]. Concerning lock-free algorithms, the use of non-preemptive commit loops was observed to result in slightly better performance. Experimental results from three representative configurations are reported in Figures 5(a)–5(c). The relevant configuration parameters are reported in the caption above each graph. Inset (a) reports results for a configuration with low contention, where lock-free algorithms exhibit slightly worse performance with respect to spin locks, but their use is not disruptive in terms of system schedulability. Insets (b) and (c) show two configurations where F|N spin locks dominate all the other mechanisms, accepting up to 65% more task sets than lock-free algorithms for $n = 30$ in inset (c), and 30% more task sets for $n = 22$ in inset (b).

B. Experiment 2: asymmetric platforms

1) *Workload generation:* To simulate asymmetric platforms, we split the available processors in two subsets: given the number of processors $m \in \{2, 4, 8\}$, we defined $\mathcal{P}^{fast} = \{P_1, \dots, P_{m/2}\}$ to be a set of “fast” processors and (ii) $\mathcal{P}^{slow} = \{P_{m/2+1}, \dots, P_m\}$ to be a set of “slow” processors. Let $s \geq 1$

be a scale factor to simulate the ratio between the different speeds of the processors. For each processor $P \in \mathcal{P}^{fast}$, we generated a task set running on P by means of the Emberson et al.’s [30] task set generator. The generator was given a target utilization U^{fast} chosen uniformly at random from the interval $[0.95, 0.99]$, and a target task count of $n' \in \{5, 7, 10\}$ tasks. For each task $T_i \in \tau(P)$, the maximum critical section (or commit loop) lengths $L_{i,q}$ were chosen uniformly at random from either $[1\mu s, 5\mu s]$ (*extremely short*), $[1\mu s, 10\mu s]$ (*very short*) or $[1\mu s, 25\mu s]$ (*short*). The same generation strategy was then applied to each processor in the set \mathcal{P}^{slow} but with a target utilization of $U^{slow} = \{0.7/s, 0.8/s\}$. Finally, the execution cost and critical section length parameters of each task executing on “slow” processors was scaled up by the factor s to coarsely simulate the effects of a slower processor. The rest of the parameters were generated as reported in Sec. VII-A1.

2) *Results*: More than 3000 different configurations have been evaluated while varying the scale factor $s \in \{1, 2, \dots, 10\}$. For each value of s , 500 different task sets were generated. Taking into account all tested configurations, we observed the following trends: (i) lock-free algorithms show significantly better performance than spin locks as the scale factor increases in case of low or moderate contention; (ii) the gap in performance between lock-free algorithms and spin locks tends to decrease as contention increases; (iii) our new analysis of F|N spin locks shows a consistent improvement over the baseline analysis [10]; (iv) F|N spin locks generally perform better than F|P spin locks; and (v) executing commit loops non-preemptively, rather than preemptively, tends to increase schedulability. The results from three representative configurations are reported in Figures 5(e), 5(d) and 5(f). As it is apparent in the graphs, all mechanisms tend to degrade in schedulability as the scale factor increases. Trend (i) can be observed in both Figures 5(e) and 5(d), where the analysis of lock-free algorithms accepts up to seven times more task sets than F|N spin locks. Trend (ii) can be observed in Figure 5(f), which reports the results for a configuration with higher contention ($p^{acc} = 0.25$). Trends (iii) and (iv) are apparent in all the three figures. Finally trend (v) is evident in Figure 5(f), where lock-free algorithms with non-preemptive commit loops are able to guarantee 10% more task sets than the preemptive alternative.

VIII. RELATED WORK

Lightweight synchronization mechanisms for multiprocessor systems have been studied for many years and comprehensive surveys are available [31, 32]. Most relevant to us, Mellor-Crummey and Scott [18] provided foundational algorithms for implementing efficient FIFO spin locks. Different techniques for implementing queue-based spin locks that allow busy-waiting jobs to be preempted were developed subsequently [28, 29]. Works addressing the efficient implementation of lock-free algorithms are also available [33, 34]. In the context of real-time systems, lock-free algorithms were first analyzed by Anderson et al. [26] under both EDF and FP scheduling on uniprocessor systems. Later, Holman and Anderson studied the use of lock-free synchronization in Pfair-scheduled real-time systems [35].

A different look at lock-free algorithms is due to Cho et al. [36], who studied workloads characterized by arrival curves managed by a uniprocessor utility-based scheduler.

Concerning spin locks, Gai et al. [10] were the first to formally analyze blocking due to spin locks, proposing the MSRP, which combines the classical SRP for uniprocessor with FIFO non-preemptive spin locks. Devi et al. [37] later extended Gai et al.’s analysis for FIFO non-preemptive spin locks to deal with globally scheduled systems. These analyses have been also used for the *Flexible Multiprocessor Locking Protocol* (FMLP) [38], which integrates FIFO non-preemptive spin locks to manage short critical sections. Finally, in 2013, Wieder and Brandenburg [12] presented the first inflation-free analysis for spin locks under P-FP scheduling; we have transferred this approach to P-EDF scheduling and adapted it to lock-free synchronization. Spin-based locks have also been investigated in work targeting reservation-based scheduling [11, 39].

Beside lightweight synchronization mechanisms, much effort has been spent on the design and the analysis of semaphore-based protocols for multiprocessor real-time systems; recent overviews are available in [1, 13, 14].

The analysis framework for P-EDF proposed in this paper relies on Baruah’s enhanced PDC [19]. In 2009, Zhang and Burns proposed the QPA algorithm [4] to speed up PDC analysis. Conceptually, we believe that our analysis framework could also benefit from integrating the QPA idea; we leave it as a future work to explore the applicability of QPA in greater detail.

IX. CONCLUSIONS AND FUTURE WORK

In this work, we have revisited the analysis of lightweight synchronization mechanisms under P-EDF scheduling.

To enable the inflation-free analysis of synchronization delays, we have developed a new PDC-based, iterative schedulability test (Sec. III). On the basis of this foundation, we have studied four different lightweight synchronization mechanisms, which for the most part had not been considered before under P-EDF, and proposed matching bounds on worst-case synchronization delay. In particular, we have proposed the first *inflation-free* analysis of lock-free algorithms under P-EDF scheduling and a new and improved analysis for FIFO spin locks.

To evaluate the proposed analyses, we conducted a large-scale experimental study considering both symmetric and asymmetric multiprocessor platforms. We observed that systems composed of asymmetric multiprocessors can benefit from the use of lock-free algorithms, especially if contention is predominantly low. At the same time, we confirmed the effectiveness of non-preemptive FIFO spin locks on symmetric multiprocessors.

The closer look at lock-free algorithms provided in this work allows for interesting future exploration. For instance, lock-free synchronization could be a good match in the domain of *component-based software design*, where it could enable the analysis of opaque components without any assumptions on maximum critical section lengths, while also isolating components from overruns during accesses to shared resources, which is the *Achilles Heel* of current component-based systems [11, 39] and a cause of much pessimism.

REFERENCES

- [1] B. Brandenburg, "Scheduling and locking in multiprocessor real-time operating systems," in *Ph.D. dissertation, The University of North Carolina at Chapel Hill*, 2011.
- [2] B. Brandenburg, J. Calandrino, and J. Anderson, "On the scalability of real-time scheduling algorithms on multicore platforms: A case study," in *Proc. of RTSS'08*, 2008.
- [3] S. Baruah, A. K. Mok, and L. E. Rosier, "Preemptively scheduling hard-real-time sporadic tasks on one processor," in *Proc. of RTSS'90*, Orlando, FL, 1990, pp. 182–190.
- [4] F. Zhang and A. Burns, "Schedulability analysis for real-time systems with EDF scheduling," *IEEE Transactions on Computers*, vol. 58, 2009.
- [5] N. Guan and W. Yi, "General and efficient response time analysis for EDF scheduling," in *Proc. of DATE'14*, 2014.
- [6] A. Bastoni, B. Brandenburg, and J. Anderson, "An empirical comparison of global, partitioned, and clustered multiprocessor EDF schedulers," in *Proc. of RTSS'10*, 2010.
- [7] B. Brandenburg, J. Calandrino, A. Block, H. Leontyev, and J. Anderson, "Real-time synchronization on multiprocessors: To block or not to block, to suspend or spin?" in *Proc. of RTAS'08*, 2008.
- [8] P. Gai, G. Lipari, L. Abeni, M. di Natale, and E. Bini, "Architecture for a portable open source real-time kernel environment," in *Proceedings of the Second Real-Time Linux Workshop and Hand's on Real-Time Linux Tutorial*, November 2000.
- [9] D. Faggioli, F. Checconi, M. Trimarchi, and C. Scordino, "An EDF scheduling class for the linux kernel," in *Proc. of RTLWS'09*, Dresden, Germany, September 28-30, 2009.
- [10] P. Gai, G. Lipari, and M. D. Natale, "Minimizing memory utilization of real-time task sets in single and multiprocessor systems-on-a-chip," in *Proc. of RTSS'01*, 2001.
- [11] A. Biondi, G. Buttazzo, and M. Bertogna, "Supporting component-based development in partitioned multiprocessor real-time systems," in *Proc. of ECRTS'15*, 2015.
- [12] A. Wieder and B. Brandenburg, "On spin locks in AUTOSAR: Blocking analysis of FIFO, unordered, and priority-ordered spin locks," in *Proc. of RTSS'13*, 2013.
- [13] B. Brandenburg, "Improved analysis and evaluation of real-time semaphore protocols for P-FP scheduling," in *Proc. of RTAS'13*, 2013.
- [14] M. Yang, A. Wieder, and B. Brandenburg, "Global real-time semaphore protocols: A survey, unified analysis, and comparison," in *Proc. of RTSS'15*, 2015.
- [15] A. Biondi and B. Brandenburg, "Lightweight real-time synchronization under P-EDF on symmetric and asymmetric multiprocessors (full version)," MPI-SWS, Tech. Rep. MPI-SWS-2016-003, May 2016, available at: <http://www.mpi-sws.org/~bbb/papers>.
- [16] M. P. Herlihy and J. M. Wing, "Linearizability: A correctness condition for concurrent objects," *ACM Transactions on Programming Languages and Systems (TOPLAS)*, vol. 12, no. 3, pp. 463–492, 1990.
- [17] T. P. Baker, "Stack-based scheduling for realtime processes," *Real-Time Systems*, vol. 3, no. 1, pp. 67–99, April 1991.
- [18] J. M. Mellor-Crummey and M. L. Scott, "Scalable reader-writer synchronization for shared-memory multiprocessors," in *Proc. of PPOPP'91*, 1991.
- [19] S. Baruah, "Resource sharing in EDF-scheduled systems: a closer look," in *Proc. of RTSS'06*, 2006.
- [20] S. Baruah, L. Rosier, and R. Howell, "Algorithms and complexity concerning the preemptive scheduling of periodic, real-time tasks on one processor," *Journal of Real-Time Systems*, vol. 2, 1990.
- [21] N. Audsley, A. Burns, M. Richardson, K. Tindell, and A. Wellings, "Applying new scheduling theory to static priority pre-emptive scheduling," *Software Engineering Journal*, vol. 8, no. 5, pp. 284–292, September 1993.
- [22] L. Thiele, S. Chakraborty, and M. Naedele, "Real-time calculus for scheduling hard real-time systems," in *In Proc. of ISCAS'00*, May 2000.
- [23] R. Pellizzoni and G. Lipari, "Feasibility analysis of real-time periodic tasks with offsets," *Real-Time Systems*, vol. 30, no. 1, pp. 105–128, 2005.
- [24] M. Spuri, "Analysis of deadline schedule real-time systems," *Technical report 2772, INRIA*, 1996.
- [25] A. C. I. Ripoll and A. Mok, "Improvement in feasibility testing for real-time tasks," *Journal of Real-Time Systems*, vol. 11, no. 1, pp. 19–39, 1996.
- [26] J. Anderson, S. Ramamurthy, and K. Jeffay, "Real-time computing with lock-free shared objects," *ACM Transactions on Computer Systems (TOCS)*, vol. 15, pp. 134–165, May 1997.
- [27] "SchedCAT: Schedulability test collection and toolkit," web site, <http://www.mpi-sws.org/~bbb/projects/schedcat>.
- [28] L. I. Kontothanassis, R. W. Wisniewski, and M. L. Scott, "Scheduler-conscious synchronization," *ACM Transactions on Computer Systems*, vol. 15, no. 1, 1997.
- [29] T. Craig, "Queuing spin lock algorithms to support timing predictability," in *In Proc. of RTSS'93*, 1993.
- [30] P. Emberson, R. Stafford, and R. I. Davis, "Techniques for the synthesis of multiprocessor tasksets," in *Proc. WATERS*, 2010.
- [31] M. Raynal, *Algorithms for mutual exclusion*. MIT Press, 1986.
- [32] J. H. Anderson, Y.-J. Kim, and T. Herman, "Shared-memory mutual exclusion: major research trends since 1986," *Distributed Computing*, vol. 16, no. 2-3, pp. 75–110, 2003.
- [33] J. Anderson and M. Moir, "Universal constructions for large objects," *IEEE Transactions on Parallel and Distributed Systems*, vol. 10, no. 12, 1999.
- [34] M. Michael and M. Scott, "Simple, fast, and practical non-blocking and blocking concurrent queue algorithms," in *In Proc. of PODC'96*, 1996.
- [35] P. Holman and J. Anderson, "Supporting lock-free synchronization in pfair-scheduled systems," *Journal of Parallel and Distributed Computing*, vol. 66, no. 1, 2006.
- [36] H. Cho, B. Ravindran, and E. D. Jensen, "Lock-free synchronization for dynamic embedded real-time systems," *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 9, 2010.
- [37] U. Devi, H. Leontyev, and J. Anderson, "Efficient synchronization under global EDF scheduling on multiprocessors," in *In Proc. of ECRTS'06*, 2006.
- [38] A. Block, H. Leontyev, B. Brandenburg, and J. Anderson, "A flexible real-time locking protocol for multiprocessors," in *In Proc. of RTCSA'07*, 2007.
- [39] D. Faggioli, G. Lipari, and T. Cucinotta, "Analysis and implementation of the multiprocessor bandwidth inheritance protocol," *Real-Time Systems*, vol. 48, no. 6, pp. 789–825, 2012.