

A Fully Preemptive Multiprocessor Semaphore Protocol for Latency-Sensitive Real-Time Applications

Björn B. Brandenburg

Max Planck Institute for Software Systems (MPI-SWS)

Abstract

Independence preservation, a property in real-time locking protocols that isolates latency-sensitive tasks from delays due to unrelated critical sections, is identified, formalized, and studied in detail. The key to independence preservation is to ensure that tasks remain fully preemptive at all times. For example, on uniprocessors, the classic priority inheritance protocol is independence-preserving. It is shown that, on multiprocessors, independence preservation is impossible if job migrations are disallowed. The $O(m)$ independence-preserving protocol (OMIP), a new, asymptotically optimal binary semaphore protocol based on migratory priority inheritance, is proposed and analyzed. The OMIP is the first independence-preserving, real-time, suspension-based locking protocol for clustered job-level fixed-priority scheduling. It is shown to benefit latency-sensitive workloads, both analytically by means of schedulability experiments, and empirically using response-time measurements in LITMUS^{RT}.

1 Introduction

A frequent, but challenging requirement in the design and implementation of practical real-time systems is the need for predictable, mutually exclusive access to *shared resources* such as I/O ports, network links, message buffers, or other shared data structures. To this end, virtually all modern real-time operating systems (RTOSs) offer some sort of locking primitives, most commonly *binary semaphores* (or *mutexes*), under which tasks suspend when trying to access already locked resources. The defining characteristics of a *real-time semaphore protocol* are that it **(i)** must guarantee bounded delays for *all* real-time tasks—which at times requires lock-holding tasks to be scheduled despite the presence of (otherwise) higher-priority tasks [30, 33]—while **(ii)** simultaneously ensuring that urgent, high-priority tasks are not unduly delayed by unrelated critical sections in lower-priority tasks.

In the *uniprocessor* case, the real-time locking problem has long been solved: the three classic uniprocessor real-time semaphore protocols, namely the *priority-inheritance protocol* (PIP) [30, 33], the *priority-ceiling protocol* (PCP) [30, 33], and the *stack resource policy* (SRP) [3], fully satisfy requirements (i) and (ii). Unfortunately, the same cannot be said of the *multiprocessor* case. While numerous protocols satisfying requirement (i) have been developed for various scheduling approaches (reviewed below), we show in this paper that *all* existing protocols fail to meet the second

requirement when confronted with *latency-sensitive* workloads (*i.e.*, if some tasks have near-zero tolerance for scheduling delays) under *partitioned scheduling* (where each task is statically assigned to a processor and each processor is scheduled individually). In other words, there exist demanding, but common hard real-time workloads that currently simply cannot be provisioned on multiprocessors under partitioned scheduling (which is widespread in practice) due to the lack of appropriate synchronization primitives.

The primary contribution of this paper is a solution to this problem: we present the design, analysis, implementation, and evaluation of the first multiprocessor real-time semaphore protocol applicable to partitioned scheduling (and other policies) that guarantees that latency-sensitive applications are *never* delayed by unrelated lower-priority tasks.

The problem. To understand why existing multiprocessor semaphore protocols fail, first consider why the classic uniprocessor protocols *do* work. For example, consider the *de facto* industry standard of preemptive *fixed-priority* (FP) scheduling in conjunction with the PIP, and let T_h denote the highest-priority task in the system. *Priority inheritance* [30, 33] famously addresses requirement (i) by temporarily raising the priority of a lock-holding task to that of the highest-priority waiter (if any). One major advantage of this approach is that task T_h is *never* delayed by critical sections associated with resources shared *only* among lower-priority tasks, since such critical sections always execute at a lower priority. *As a result, such unrelated critical sections may be of arbitrary length without affecting the temporal correctness of T_h .* This isolation property, which we formalize as “independence preservation” in Sec. 3, crucially depends on the fact that lock-holding tasks remain fully preemptive.

Now consider the most straightforward (and widespread) multiprocessor extension of FP scheduling, namely *partitioned fixed-priority* (P-FP) scheduling, together with the *multiprocessor priority-ceiling protocol* (MPCP) [29, 30], a classic shared-memory semaphore protocol for P-FP scheduling. Under partitioned scheduling, it is particularly difficult to ensure that lock-holders are scheduled when remote tasks are blocked;¹ the MPCP—and all real-time semaphore protocols for partitioned scheduling developed since—employ *priority boosting*, that is, they simply raise the priority of lock-holding tasks above any “normal” priority, thereby ensuring that lock-holding tasks are *never*

¹Priority inheritance is ineffective under partitioning, *e.g.*, see [8, 12].

preempted by non-lock-holding tasks. While effective, this simple technique has the unfortunate consequence that lock-holding tasks are effectively no longer fully preemptive: even the highest-priority task T_h cannot preempt all other tasks at all times. *As a result, there is no isolation: the temporal correctness of any task depends on the maximum critical section length of any other task (assigned to the same processor).* Crucially, this is unacceptable for latency-sensitive workloads, where lower-priority tasks must *not* affect the correctness of the highest-priority tasks.

The described phenomenon is not specific to the MPCP, but rather a consequence of priority boosting. While priority boosting can be acceptable if *all* critical sections can be ascertained with high confidence to always be “short enough,” latency-sensitive workloads with highly heterogeneous timing requirements fundamentally require high-priority tasks to be shielded from unrelated critical sections. Such workloads are common in practice. For instance, in the automotive domain, it is not unusual for real-time constraints to span three orders of magnitude, ranging from a few milliseconds for high-frequency control tasks up to one second for less urgent tasks (*e.g.*, see [36]). However, in a task with a period of 1000 milliseconds, it is not unreasonable to encounter critical sections with *worst-case* lengths in excess of one millisecond, especially in embedded systems with relatively slow processors. Clearly such a critical section cannot be permitted to delay latency-sensitive tasks (*e.g.*, those with a one-millisecond period). However, as we discuss next, *none* of the real-time semaphore protocols for partitioned scheduling proposed to date can actually guarantee this—a major oversight that we rectify in this paper.

Related work. Multiprocessor semaphore protocols are typically designed for, and tightly integrated with, a specific scheduling policy. Most relevant to our work are prior semaphore protocols for partitioned scheduling. The first such protocols were proposed by Rajkumar *et al.* [29–31], who designed the aforementioned MPCP and the *distributed PCP* for P-FP-scheduled systems. In later work, Block *et al.* [7] proposed the *flexible multiprocessor real-time locking protocol* (FMLP), and an improved version for partitioned scheduling, termed FMLP⁺, has since been developed in [8]. In recent work, Nemati *et al.* [27] proposed a semaphore protocol that simplifies the consolidation of independently developed (legacy) applications onto P-FP-scheduled multicore platforms. Crucially, all of the just-cited protocols employ priority boosting and are thus not suitable for latency-sensitive workloads, as explained above.

In work focused on *global scheduling*, under which (conceptually) all processors serve a single ready queue and tasks may migrate freely among cores, several semaphore protocols have been proposed in recent years, including a global variant of the FMLP and two extensions of the PCP [20, 26]. These protocols rely on priority inheritance [33] instead of priority boosting and allow tasks to remain fully preemptive;

however, they apply *only* to global scheduling and cannot be easily transferred to the partitioned case, which is (more) commonly used in industrial applications (*e.g.*, partitioning is mandated by the Autosar 4.0 standard).

Most closely related to this paper are optimality results [8, 11, 12] pertaining to the intuitive goal of “minimal blocking” in multiprocessor real-time semaphore protocols under *job-level fixed-priority* (JLFP) scheduling, a broad category of schedulers that includes the FP and *earliest-deadline first* (EDF) policies. Specifically, recent work [11] investigated asymptotic bounds on maximum *priority-inversion blocking* (*pi-blocking*), which (intuitively) occurs whenever a high-priority task is delayed by a lower-priority one. Surprisingly, there exist two classes of schedulability analysis, called *suspension-aware* and *suspension-oblivious* analysis [11], resp., that give rise to different lower bounds on worst-case pi-blocking, namely $\Omega(n)$ and $\Omega(m)$, resp., where n denotes the number of tasks and m the number of processors. In brief, the difference arises due to how semaphore-induced self-suspensions are accounted for.

This paper pertains to the suspension-oblivious case, for which there already exists the asymptotically optimal family of $O(m)$ *locking protocols* (OMLP) [8, 11, 12]. In particular, the OMLP family contains semaphore-based mutual exclusion protocols for global scheduling, partitioned scheduling, and for *clustered scheduling* [4, 16], which is a hybrid (or generalization) of global and partitioned scheduling where tasks are statically assigned to non-overlapping subsets (or *clusters*) of processors, and each cluster is scheduled using a global policy. We revisit these protocols, denoted as P-OMLP, G-OMLP, and C-OMLP, resp., in greater detail in Sec. 2. However, note that the OMLP family decidedly does *not* solve the problem addressed in this paper, which is to support *latency-sensitive* workloads under *clustered* scheduling, since the priority-inheritance-based G-OMLP applies *only* to global scheduling, whereas the P-OMLP and the C-OMLP use (a variant of) priority boosting, and thus are susceptible to potentially problematic latency increases.

A well-studied, sometimes-preferable alternative to semaphores are *non-preemptive spinlocks*, in which blocked jobs busy-wait by executing a delay loop (*e.g.*, see [7, 8, 17, 24]). While busy-waiting conceptually wastes processor cycles, non-preemptive spinlocks are an attractive choice in practice due to their simplicity and low overheads [8, 14]. Unfortunately, non-preemptive spinlocks can cause unacceptable delays for latency-sensitive workloads because tasks may remain non-preemptive for prolonged times. While it is possible to reduce the negative effects by allowing spinning tasks to remain preemptive (*e.g.*, see [2, 19, 34]), critical sections are still executed non-preemptively in such spinlocks and thus have a major impact on latency, not unlike the priority-boosting of critical sections under the MPCP.

The need for non-preemptive execution can be avoided altogether with *helping* mechanisms (*e.g.*, see [35]), where instead of executing a delay loop, “spinning” tasks attempt

to complete the critical sections of preempted lock holders. However, such helping-based approaches have two limitations in practice. First, they are considerably more complicated than regular spinlocks, which negates one of the key reasons to use spinlocks in the first place, and second, they are less general than regular locks. That is, “helping” cannot be offered as a generic OS primitive, but must be integrated individually into each application because critical sections must be implemented in a “helping-aware” manner such that they can be posted to a queue of pending operations.

Finally, in work aimed at temporal isolation among tasks in mixed real-time/non-real-time environments, Faggioli *et al.* [23] presented the *multiprocessor bandwidth-inheritance protocol* (MBWI), a scheduler-agnostic locking protocol based on preemptive spinning. Notably, to ensure progress, the MBWI allows lock-holding tasks to migrate to processors on which tasks are waiting for the lock, an effective technique that we reuse in Sec. 3.

Contributions. We provide a precise definition of *independence preservation*, which captures the idea that “tasks should be isolated from unrelated critical sections,” and provide evidence of its considerable practical importance (Sec. 3). We show in Sec. 4.1 that this desirable property is impossible on multiprocessors unless jobs may (temporarily) migrate among all processors. Based on this observation, we present in Sec. 4.2 our main contribution, the $O(m)$ *independence-preserving locking protocol* (OMIP) for clustered scheduling, and show that it is optimal within a factor of approximately two under suspension-oblivious analysis (Sec. 4.3). Finally, in Sec. 5, we report on response-time measurements and schedulability experiments that show the OMIP to greatly benefit latency-sensitive workloads. We begin with a review of needed background.

2 Background and Definitions

This paper assumes the classic sporadic task model and follows the conventions from [11].

Tasks and jobs. The workload is specified as a set of n sporadic tasks $\tau = \{T_1, \dots, T_n\}$ that is to be scheduled on m identical processors. We let T_i denote a task with a *worst-case per-job execution time* e_i and a *minimum job separation* p_i , and let J_i denote a job of T_i . Our results do not depend on the type of deadline constraint; we assume implicit deadlines for simplicity. A task’s *utilization* is the ratio $u_i = e_i/p_i$, and τ ’s *total utilization* is given by $U = \sum_{T_i \in \tau} u_i$.

A task’s *slack* $p_i - e_i$ describes how much delay one of its jobs can tolerate without missing its deadline. A task is *latency-sensitive* if its slack does not exceed the maximum critical section length of other tasks (see below).

Jobs are *pending* from the time when they are released until they complete. A pending job can be in one of two states: a *ready* job is available for execution, whereas a *suspended* job cannot be scheduled. A job *resumes* when its state changes from suspended to ready.

Scheduling. Under *clustered scheduling* [4, 16], processors are grouped into $\frac{m}{c}$ non-overlapping sets (or *clusters*) of c processors each, where C_k denotes the k^{th} cluster. We require uniform cluster sizes and for simplicity assume that m is an integer multiple of c .² Each task is statically assigned to a cluster; we let $C(T_i)$ denote the cluster to which task T_i has been assigned. *Partitioned* ($c = 1$) and *global* ($c = m$) scheduling are special cases of clustered scheduling.

Priority. Each cluster is scheduled independently using a work-conserving *job-level fixed-priority* (JLFP) scheduler. Under a JLFP scheduler, each job is assigned a unique *base priority* that does not change while it is pending. However, a locking protocol (see below) may temporarily raise the *effective priority* of a job to ensure its timely completion. At any point in time and w.r.t. each cluster, the c highest-effective-priority ready jobs are scheduled.

JLFP scheduling is a broad category of scheduling policies that includes EDF and FP scheduling. Note that in the following, we use the terms “higher-priority” and “lower-priority” to refer to the generic JLFP notion of a job’s base priority (and *not* to a task’s priority under FP scheduling). For illustrative purposes, we consider *clustered earliest-deadline-first* (C-EDF) scheduling as a representative JLFP policy. (We chose C-EDF since it is the policy actually implemented in LITMUS^{RT}, the RTOS underlying the evaluation in Sec. 5.) We let P-EDF denote the special case of partitioned scheduling (*i.e.*, C-EDF scheduling with $c = 1$).

Real-time locking. Besides the m processors, tasks share r serially-reusable shared resources ℓ_1, \dots, ℓ_r . Access to shared resources is governed by a *locking protocol* that ensures mutual exclusion by serializing conflicting requests.

A job J_i *requests* a resource ℓ_q at most $N_{i,q}$ times. If ℓ_q is already in use, J_i incurs *acquisition delay* until its request for ℓ_q is *satisfied*. In this paper, we focus on semaphore protocols, that is, jobs wait by self-suspending. Once J_i has finished using ℓ_q , its request is *complete* and J_i *releases* ℓ_q . We let $L_{i,q}$ denote the *maximum critical section length*, that is, the maximum time that J_i uses ℓ_q as part of a single request (or *critical section*). We assume that J_i must be scheduled in order to use ℓ_q ; however, the analysis in Sec. 4 remains valid even if critical sections contain suspensions (*e.g.*, when performing I/O). For notational convenience, we define $L_q^{\max} \triangleq \max_i \{L_{i,q}\}$ and $L^{\max} \triangleq \max_q \{L_q^{\max}\}$, and require that $L_{i,q} = 0$ if $N_{i,q} = 0$.

We assume that tasks do not hold resources across job boundaries and that jobs request at most one resource at any time. Analysis of nested critical sections is beyond the scope of this paper and the subject of future work (see Sec. 6).

Priority inversion. Intuitively, a *priority inversion* exists if a high-priority job that *should* be scheduled is *not* scheduled, which is typically due to acquisition delay. However,

² If c does not evenly divide m , there are $\lfloor \frac{m}{c} \rfloor$ clusters of size c and one cluster of size $m \bmod c$. This does not affect the asymptotic bounds.

priority inversion is not the same as acquisition delay: a suspended job does not suffer a priority inversion if it would not have been scheduled anyway (*i.e.*, in the presence of higher-priority jobs) [30, 33]. We let b_i denote the maximum cumulative duration of priority inversion incurred by any J_i .

Recent work [8, 11] showed that what exactly constitutes a priority inversion (Def. 1 below) depends on how self-suspensions are accounted for during schedulability analysis (*i.e.*, when checking whether all jobs will meet their deadline). Accurate analysis of self-suspensions is NP-hard [32] and complicates schedulability tests considerably. Therefore, many published tests (*e.g.*, [5, 6]) make the simplifying assumption that tasks never self-suspend. Such *suspension-oblivious* (*s-oblivious*) tests can still be used with semaphores by inflating each task’s execution cost e_i by b_i prior to applying the test [11]. In contrast, *suspension-aware* (*s-aware*) schedulability tests explicitly account for self-suspensions. The focus of this paper is s-oblivious analysis; the s-aware case is subject of ongoing work. S-oblivious analysis yields the following definition of “priority inversion.”

Def. 1. A pending job J_i suffers an *s-oblivious priority inversion* at time t iff J_i is (i) not scheduled and (ii) among the c highest-base-priority pending jobs in cluster $C(T_i)$ [11].

Note that clause (ii) pertains only to jobs in J_i ’s assigned cluster $C(T_i)$ —the priorities of jobs in other clusters are irrelevant since each cluster is analyzed individually—and that it pertains to pending (and not just ready) jobs because self-suspensions are analyzed as execution time.

Priority inversions are considered to cause “blocking” because they lead to an undesirable increase in worst-case response time. To avoid confusing such locking-related delays with other uses of the word “blocking,” we use the more specific term “pi-blocking” in this paper.

PI-Blocking optimality. The purpose of a real-time locking protocol is to minimize the occurrence and duration of priority inversions, which, however, cannot be avoided entirely. The significance of differentiating between s-oblivious and s-aware analysis is that they give rise to different lower bounds on *maximum pi-blocking*, formally $\max_{\tau} \{b_i\}$ [11]. In particular, the s-oblivious case is subject to an $\Omega(m)$ lower bound on maximum pi-blocking, whereas s-aware analysis yields an $\Omega(n)$ lower bound (where $n \geq m$). In other words, under s-oblivious analysis, there exist pathological task sets such that pi-blocking is linear in the number of processors under *any* semaphore protocol. To be asymptotically optimal, a locking protocol must ensure that maximum pi-blocking is *always* within a constant factor of m . In the suspension-oblivious case, $\max_i \{b_i\} = O(m)$ is hence asymptotically optimal. When stating asymptotic bounds, L^{max} and $\sum_q N_{i,q}$ are considered constant (*i.e.*, not a function of n or m) [11, 12]. The task execution costs e_i are *not* constants, and pi-blocking is considered “bounded” only if no b_i depends on any e_i .

The OMLP. The family of $O(m)$ locking protocols is a collection of in total five asymptotically optimal locking proto-

cols for global, partitioned, and clustered JLFP scheduling [8, 11, 12]. Most relevant to this paper are the two mutual exclusion protocols for partitioned [11] and clustered scheduling [12], which we denote as P-OMLP and C-OMLP, resp. Both use *priority boosting* to ensure lock-holder progress, which means that the effective priority of lock-holding jobs exceeds any possible base priority. Lock-holding jobs thus are never preempted by non-lock-holding jobs, and in particular not by newly-released jobs (which cannot yet have acquired any locks at the time of their release).



While we consider clustered JLFP scheduling with arbitrary cluster sizes, we focus on the P-OMLP [11] as a baseline because it is somewhat simpler and because it suffices to reveal the limitations of priority boosting. The key difference between the two protocols is that the P-OMLP is asymptotically optimal only for $c = 1$, whereas the C-OMLP uses a refined version of priority boosting called *priority donation* that allows it to cover the entire range of $1 \leq c \leq m$ [12]. For the purpose of this paper, the C-OMLP and the P-OMLP can be considered to be equivalent.

In brief, the P-OMLP works as follows. Requests for each ℓ_q are serialized with a FIFO queue FQ_q . However, before a job can be enqueued in any FQ_q , it must hold its processor’s *contention token* CT_k . As there is only one such token per processor, at most m jobs are enqueued in any FQ_q at any time, which yields the desired $O(m)$ bound. Contention for each CT_k is managed using a per-processor priority queue PQ_k . The P-OMLP is simple, easy to implement [8], ensures asymptotically optimal maximum pi-blocking, and has been shown to perform well compared to the FMLP⁺ and the MPCP for certain workloads [8, 12]. However, its chief limitation, which it shares with the FMLP⁺ and the MPCP, is its reliance on priority boosting, as we show in detail next.


3 Independence Preservation

In this section, we first give a precise specification for the intuitive requirement that “high-priority jobs should be isolated from unrelated critical sections,” and then provide evidence of its considerable practical importance.


To motivate the need for such isolation, consider the example P-OMLP schedule given in Fig. 1(a), which demonstrates that priority boosting is ill-suited to latency-sensitive workloads. In this particular example, task T_1 with parameters $e_1 = 6$ and $p_1 = 10$ is latency-sensitive since it has only four time units of slack, less than task T_2 ’s maximum critical section length ($L_{2,1} = 5$). In the depicted schedule, job J_2 acquires resource ℓ_1 and becomes priority-boosted shortly after its release at time 0. Hence, job J_1 cannot preempt J_2 at time 1 and consequently misses its deadline at time 11. Notably, J_1 does not access shared resources and is continuously the highest-priority job while it is pending—clearly it *should* not incur *any* delays, but under the P-OMLP (or any other semaphore protocol based on priority boosting), J_1 is not isolated from J_2 ’s long critical section because J_2



(a) Priority boosting causes pi-blocking for independent jobs.



(b) Priority inheritance is ineffective across clusters.



(c) Temporarily migrating J_2 to C_1 renders the task set schedulable.

Figure 1: P-EDF schedules ($c = 1$) of three tasks T_1, T_2, T_3 with execution costs $e_1 = 6$, $e_2 = 9$, $e_3 = 3$, and periods $p_1 = 10$, $p_2 = 40$, and $p_3 = 11$. Tasks T_2 and T_3 share one resource (ℓ_1), where $L_{2,1} = 5$ and $L_{3,1} = 1$. While P-EDF is assumed in this example, the same effect can arise under any partitioned JLFP scheduler (including P-FP). The legend applies to Fig. 2, too.

is effectively non-preemptable while accessing ℓ_1 .

It bears repeating that classic priority inheritance [33], which allows jobs to remain fully preemptive, is unfortunately ineffective (analytically speaking) when applied *across* clusters or partitions, as shown in Fig. 1(b). At time 2, job J_3 attempts to lock ℓ_1 , but must wait since J_2 already holds the resource. Assuming priority inheritance, this causes J_2 's effective priority (*i.e.*, deadline under P-EDF) to be raised from time 40 (J_2 's deadline) to time 12 (J_3 's deadline). However, J_1 has an even earlier deadline at time 11 and J_2 remains preempted. As a result, J_3 incurs *unbounded* pi-blocking in the sense that it is no longer possible to bound pi-blocking only in terms of critical section lengths (b_3 depends on e_1). A deadline miss results.

However, a simple tweak renders this example schedulable: instead of boosting J_2 's priority while holding ℓ_1 , it should simply be migrated to J_3 's cluster C_1 when J_3 requests ℓ_1 . As illustrated in Fig. 1(c), this allows J_2 to complete its critical section in a timely fashion so that J_3

can meet its deadline. Notably, this is accomplished without causing *any* additional delay for the *independent*, latency-sensitive job J_1 , and while some pi-blocking is unavoidable when using semaphores [11], it is only incurred by the job that actually requires a shared resource (J_3 in this example). We formalize this property as follows (first defined in [8]).

Def. 2. Let $b_{i,q}$ denote the maximum pi-blocking incurred by J_i due to requests by *any* task for resource ℓ_q . Under s-oblivious analysis, a locking protocol is *independence-preserving* if and only if $N_{i,q} = 0$ implies $b_{i,q} = 0$.

In other words, independence among tasks is preserved if the overall s-oblivious³ pi-blocking $b_i = \sum_q b_{i,q}$ does not depend on resources that a task does not access.

Significance. We posit that independence preservation is indispensable for many workloads common in current practice in industry, as illustrated by the following use cases.

First, if the locks acquired (only) by lower-priority tasks protect data structures of potentially unbounded size (*e.g.*, lists or trees of unknown size), or if they serialize operations of potentially unbounded or unknown runtime (*e.g.*, vendor-provided “black-box” functionality), then it is not possible to derive *a priori* bounds on the maximum duration that lower-priority tasks remain effectively non-preemptive. This renders priority-boosting-based locking protocols unviable for real-time systems that need to be certified, even if the developer is certain that all lower-priority critical sections are “short enough.” Due to the lack of isolation under priority boosting, a small, but hard-to-quantify risk remains, which can pose considerable challenges in the certification process. Use of a predictable, independence-preserving locking protocol allows to categorically rule out such concerns.

Second, in open systems (*i.e.*, in systems in which the final task set composition is unknown at design time), in RTOSs based on commodity OSs, and in real-time applications incorporating commodity libraries (*e.g.*, for common tasks such as encryption, compression, or signal processing), locks acquired by lower-priority (or even best-effort) tasks often stem from (a large volume of) untrusted code. In this case, it is (at least economically) impossible to rule out the existence of unknown, potentially dangerous critical sections. For instance, this is a major problem in real-time Linux variants, since the Linux kernel contains millions of lines of code of varying quality (device drivers are notorious for above-average defect densities [18]), and because real-time locking primitives are exported to *all*, potentially untrusted user-space tasks (via the *pthread*s library). With an independence-preserving locking protocol, tasks remain fully preemptive and it is sufficient to inspect and test only the (few) parts of the kernel actually used by latency-sensitive tasks.

Finally, current industrial practice heavily favors fully preemptive kernel designs and uses scheduling latency (*i.e.*, how quickly a newly released high-priority job can be sched-

³Def. 2 is specific to s-oblivious analysis. For technical reasons, s-aware analysis requires a slightly weaker definition (a subject of future work).

uled) as a key evaluation metric. For example, the non-profit *Open Source Automation Development Lab* reports long-term, large-scale scheduling latency measurements of Linux-based real-time systems to demonstrate Linux’s viability as an RTOS [28]. Similarly, commercial RTOSs are commonly advertised as sporting fully preemptive designs and negligible scheduling latencies. Case in point, a prominently advertised feature of the commercial *ThreadX* uniprocessor RTOS is that “[a] high priority thread starts responding to an external event on the order of the time it takes to perform a highly optimized *ThreadX* context switch” [22]. Without an independence-preserving locking protocol, it will be impossible to extend such claims to partitioned scheduling.

Taken as a whole, these examples reveal a clear need for practical real-time multiprocessor semaphore protocols that are both *predictable* and *independence-preserving*. However, in the case of partitioned and clustered scheduling, no such protocols have been proposed to date. To address this gap, we present the first design and analysis of an independence-preserving real-time semaphore protocol for clustered scheduling in the remainder of this paper.

4 The OMIP for Clustered JLFP Scheduling

Real-time semaphore protocols must deal with two central issues: how to avoid excessive pi-blocking when lock-holding jobs are preempted (Sec. 4.1), and how to order conflicting requests such that pi-blocking is minimized to the extent possible (Sec. 4.2). We begin with the former aspect.

4.1 An Independence-Preserving Progress Mechanism

As reviewed in Sec. 1, all prior real-time semaphore protocols for partitioned and clustered scheduling use variants of priority boosting to ensure lock-holder progress, which is not independence-preserving. A new progress mechanism is thus required. Interestingly, the example schedule in Fig. 1(c) contains an *inter-cluster migration* (i.e., job J_2 must temporarily execute in cluster C_1 despite T_2 being assigned to C_2), which runs counter to the idea of clustered scheduling. Such migrations are in fact unavoidable.

Theorem 1. *Under clustered JLFP scheduling with $c \neq m$, it is impossible for a semaphore protocol to simultaneously (i) prevent unbounded pi-blocking, (ii) be independence-preserving, and (iii) avoid inter-cluster migrations.*

Proof. Consider Fig. 1. When job J_3 issues its request at time 2, any locking protocol that does *not* employ inter-cluster migrations has only two options: either it forces J_2 to complete its critical section at the expense of delaying J_1 , as shown in Fig. 1(a), or the locking protocol can isolate J_1 from any locking-related delays, as shown in Fig. 1(b). In the former case, the locking protocol is not independence-preserving, in the latter case, J_3 is subject to unbounded pi-blocking because b_3 then depends on J_1 ’s execution cost e_1 (recall that pi-blocking is considered “bounded” only if it

does not depend on job execution costs). \square

Priority boosting satisfies properties (i) and (iii) at the expense of (ii), whereas priority inheritance satisfies properties (ii) and (iii) at the expense of (i). For our purpose, however, we require a progress mechanism that satisfies both properties (i) and (ii), which according to Theorem 1 is only possible at the expense of (iii). As it turns out, such a mechanism is easily obtained by generalizing the idea of priority inheritance to “cluster inheritance,” that is, inheritance should not only pertain to scheduling priorities, but also to *where* a job is eligible to execute, a concept that we refer to as “migratory priority inheritance” [13].

Def. 3. Let W_i denote the set of jobs waiting for a job J_i to release a resource. Under *migratory priority inheritance*, whenever J_i is not scheduled (but ready) and there exists a job $J_x \in W_i \cup \{J_i\}$ such that J_x is eligible to be scheduled in its assigned cluster (i.e., whenever J_i is preempted and there are fewer than c ready higher-effective-priority jobs in J_x ’s cluster), J_i migrates to J_x ’s cluster (if necessary) and assumes J_x ’s priority. After releasing all locks, J_i migrates back to its assigned cluster (if necessary).

Paraphrased, the idea is to move preempted, lock-holding jobs among clusters such that they are always local to a waiting job that would be scheduled if it were not suspended (if any). Migratory priority inheritance may be triggered whenever a lock-holding job J_i is preempted, when a new waiter is added to W_i , or when one of the jobs in W_i becomes one of the c highest-effective-priority jobs in its cluster (due to the completion or suspension of a higher-priority job). If there are multiple candidates in $W_i \cup \{J_i\}$, then J_x may be chosen arbitrarily among the candidates (e.g., cache locality may be taken into consideration).

Migratory priority inheritance is a straightforward generalization of classic priority inheritance [33]; it can also be understood as a greatly simplified, reduced-to-the-core⁴ version of Faggioli *et al.* multiprocessor bandwidth inheritance [23]. Notably, a progress mechanism akin to migratory priority inheritance appeared as early as 2001 in TU Dresden’s Fiasco microkernel under the name “local helping” [25] (a more detailed discussion of these techniques can be found in [13]). While migratory priority inheritance has received surprisingly little attention in the real-time literature so far, it is ideal for our purposes as ready jobs are never pi-blocked since lock holders remain fully preemptive.

Lemma 1. *In a mutual-exclusion semaphore protocol based on migratory priority inheritance, if a job J_i is ready at time t , then it does not incur s-oblivious pi-blocking at time t .*

Proof. Follows from the fact that priorities are not “duplicated” in a mutual-exclusion protocol. For each job that inherits a priority higher than J_i ’s base priority in cluster $C(T_i)$, there exists at least one suspended, pending job local

⁴The MBWL [23] also uses (non-optimal) FIFO queuing and allows for budget overruns, which is orthogonal to the idea of migrating lock holders.

to $C(T_i)$ with base priority higher than J_i 's base priority. If J_i is ready but not scheduled at time t , then there exist c jobs with higher effective priorities, which, due to the one-to-one inheritance relationship (and because priorities are inherited only locally after migrating to the blocked job's cluster), implies the existence of c pending higher-base-priority jobs local to $C(T_i)$, which precludes s-oblivious pi-blocking. \square

However, migratory priority inheritance by itself is *not* sufficient to guarantee independence preservation. For example, if combined with the P-OMLP's contention token abstraction, requests of otherwise independent jobs could still conflict (*i.e.*, contention tokens create artificial dependencies). Rather, a new locking protocol designed specifically to take advantage of migratory priority inheritance is required.

4.2 Queue Structure for Optimal Waiting Times

Asymptotic optimality under s-oblivious analysis requires $\max_{\tau} b_i = O(m)$. Clearly, this cannot be achieved using only FIFO queues, which would yield $\Omega(n)$ pi-blocking, nor can it be achieved using only priority queues, which are susceptible to starvation and can be shown to yield $\Omega(m \cdot n)$ pi-blocking [12]. However, a hybrid FIFO/priority queue design *can* be used to obtain $O(m)$ pi-blocking bounds under *global* scheduling [11]. While such a hybrid queue cannot be *directly* applied to clustered scheduling (this would require comparing priorities across cluster boundaries, which, analytically, is meaningless), it can be used *within* each cluster. Further, while it is not possible to directly apply the P-OMLP's per-processor contention token design, it is beneficial to reuse the underlying idea of multiple "stages" of contention; that is, contention should be resolved first on an intra-cluster basis (where priorities can be safely compared) and only thereafter on an inter-cluster basis. Based on these considerations, the $O(m)$ *independence-preserving locking protocol* (OMIP) is defined as follows.

Structure. Each shared resource ℓ_q is protected by a global FIFO queue GQ_q of maximum length $\frac{m}{c}$. The job at the head of GQ_q holds ℓ_q . Access to GQ_q is resolved on a per-cluster and per-resource basis: in each cluster C_k , there exist another two queues for each ℓ_q , a bounded-length FIFO queue $FQ_{q,k}$ of maximum length c that feeds into GQ_q , and a priority queue $PQ_{q,k}$ that feeds into $FQ_{q,k}$.

Rules. Requests for each ℓ_q are satisfied as follows. Let J_i denote a job of a task assigned to cluster C_k . Conceptually, J_i first enters its local $PQ_{q,k}$ and then advances through the queues until it becomes the head of GQ_q .

- R1** When J_i issues a request for ℓ_q and $FQ_{q,k}$ is empty, then J_i is enqueued in both GQ_q and $FQ_{q,k}$. Otherwise, if there are fewer than c jobs queued in $FQ_{q,k}$, then J_i is enqueued only in $FQ_{q,k}$. Finally, if there are already c jobs queued in $FQ_{q,k}$, then J_i is enqueued in $PQ_{q,k}$.
- R2** J_i 's request for ℓ_q is satisfied when it becomes the head of GQ_q . J_i is suspended while it waits (if necessary).




Figure 2: C-EDF schedule of six jobs in two two-processor clusters sharing one resource ℓ_1 under the OMIP. Changes in shade indicate migrations; jobs J_3 and J_6 migrate across cluster boundaries due to migratory priority inheritance. See Fig. 1 for a legend.

- R3** While J_i holds ℓ_q , it benefits from migratory priority inheritance (w.r.t. any job waiting to acquire ℓ_q).
- R4** When J_i releases ℓ_q , it is dequeued from both GQ_q and $FQ_{q,k}$. If $PQ_{q,k}$ is non-empty, then the head of $PQ_{q,k}$ is transferred to $FQ_{q,k}$. Further, if $FQ_{q,k}$ is non-empty, then the new head of $FQ_{q,k}$ is enqueued in GQ_q . The new head of GQ_q , if any, is resumed.

Rules R1–R4 together ensure $O(m)$ s-oblivious pi-blocking, which we show in Sec. 4.3 below after a brief example.

Example. Fig. 2 shows an example OMIP schedule under C-EDF scheduling with $m = 4$ and $c = 2$. Job J_3 enters both GQ_1 and $FQ_{1,1}$ and acquires ℓ_1 at time 1 since ℓ_1 is uncontested (Rule R1). Job J_6 requests ℓ_1 shortly thereafter and enqueues in GQ_1 and $FQ_{1,2}$ (Rule R1), but must suspend since J_3 holds ℓ_1 (Rule R2). At time 2, J_3 is preempted due to the release of J_1 and J_2 . As J_3 holds ℓ_1 , it is subject to migratory priority inheritance and migrates to J_6 's cluster, where it is scheduled on processor 3 (Rule R3). J_2 proceeds to request ℓ_1 as well and enqueues in $FQ_{1,1}$, but suspends without enqueueing in GQ_1 because $FQ_{1,1}$ already contained J_3 (Rule R1). J_3 is now eligible to execute in C_1 again, but remains scheduled in C_2 since Def. 3 only takes effect when J_3 is preempted. At time 4, jobs J_4 and J_5 are released in C_2 , which implies that J_6 is no longer among the $c = 2$ highest-priority pending jobs. J_3 is thus preempted and migrates back to C_1 (Def. 3), where it inherits J_2 's priority (Rule R3). When J_3 releases ℓ_1 at time 5, it is removed from both GQ_1 and $FQ_{1,1}$, and the new head of $FQ_{1,1}$, which is J_2 , is enqueued in GQ_1 (Rule R4). J_2 does not immediately acquire ℓ_1 because it is preceded by J_6 (Rule 2). However, due to the presence of J_4 and J_5 , J_6 is no longer of sufficient priority to be scheduled in C_2 . Hence, migratory priority inheritance takes effect (Rule 3) and J_6 migrates to C_1 to execute its critical section. At time 6, J_6 releases ℓ_1 and ceases to benefit from migratory priority inheritance; it is thus preempted and not scheduled until time 7 when J_5 completes. J_2 finally acquires ℓ_1 as it becomes the head of

GQ₁. Importantly, the example shows that jobs J_1 , J_4 , and J_5 , which do not require ℓ_1 , do not incur any pi-blocking.

4.3 PI-Blocking Optimality

We bound the maximum s-oblivious pi-blocking incurred by a given task T_i under the OMIP. In the following, let J_i denote a job in cluster C_k that has requested resource ℓ_q . We begin with the observation that migratory priority inheritance ensures lock-holder progress.

Lemma 2. *Let J_h denote the job at the head of GQ_q (i.e., the job that holds ℓ_q). If J_i incurs s-oblivious pi-blocking at time t , and if J_h requires processor service at time t (i.e., if J_h is ready), then J_h is scheduled.*

Proof. Since J_i incurs s-oblivious pi-blocking, it is among the c highest-priority jobs in C_k . Therefore, J_i would be eligible to execute if it were not suspended. By Rule R3, J_h is subject to migratory priority inheritance. Thus, by Def. 3, if J_h requires processor service to proceed (i.e., if J_h is ready), then J_h migrates to $C(T_i)$ (if required), inherits J_i 's priority, and is thus scheduled by a processor in $C(T_i)$. \square

Lemma 2 implies that transitive pi-blocking due to processor unavailability is impossible, which allows us to bound maximum pi-blocking in terms of the number of critical sections that must complete before J_i acquires ℓ_q . Due to the multi-stage structure of the OMIP, we proceed in reverse order, starting with jobs already enqueued in GQ_q.

Lemma 3. *Once J_i is enqueued in GQ_q, at most $\frac{m}{c} - 1$ requests for ℓ_q must complete before J_i holds ℓ_q .*

Proof. Due to the interplay of Rules R1 and R4, at most one job per cluster is enqueued in GQ_q at any time. Thus, as there are $\frac{m}{c}$ clusters and since GQ_q is FIFO-ordered, at most $\frac{m}{c} - 1$ critical sections must complete before J_i becomes the head of the queue and its request is satisfied (Rule 2). \square

This also bounds how long J_i remains the head of FQ_{q,k}.

Lemma 4. *Once J_i becomes the head of FQ_{q,k}, at most $\frac{m}{c}$ requests for ℓ_q must complete before J_i is dequeued.*

Proof. Follows from Lemma 3 and the fact that J_i must complete its own request before being dequeued. \square

This yields a bound on how long it takes to traverse FQ_{q,k}.

Lemma 5. *Once J_i is enqueued in FQ_{q,k}, at most $\frac{m}{c} \cdot (c - 1)$ requests for ℓ_q must complete before J_i enters GQ_q.*

Proof. J_i enters GQ_q when it becomes the head of FQ_{q,k} (Rules R1 and R4). Due to the bounded length of FQ_{q,k}, at most $c - 1$ jobs precede J_i in FQ_{q,k}. By Lemma 4, each such job, once it has become the head of FQ_{q,k} is dequeued after at most $\frac{m}{c}$ requests complete. Thus there are no jobs remaining in front of J_i after $\frac{m}{c} \cdot (c - 1)$ requests complete. \square

The final step is to bound the pi-blocking that J_i incurs while waiting in PQ_{q,k}. Since priority queues do not offer strong progress guarantees, simply counting preceding jobs is insufficient in this case and an indirect proof is required.

Lemma 6. *While J_i waits in PQ_{q,k}, at most m requests for ℓ_q must complete until either J_i enters FQ_{q,k} or J_i stops to incur s-oblivious pi-blocking.*

Proof. Suppose not. Then there exists a point in time t at which more than m requests have completed since J_i entered PQ_{q,k} and J_i still waits in PQ_{q,k} and incurs pi-blocking. By Lemma 4, a job at the head of FQ_{q,k} is dequeued after at most $\frac{m}{c}$ requests complete. Thus, at time t , at least $m/\frac{m}{c} = c$ jobs have been dequeued from FQ_{q,k} since J_i entered PQ_{q,k}. By Rule R4, the highest-priority job in PQ_{q,k} is transferred to FQ_{q,k} when the head of FQ_{q,k} is dequeued. Since J_i is still in PQ_{q,k} at time t , a job with higher priority has been moved to FQ_{q,k} at least c times. Therefore, each job in FQ_{q,k} has a higher priority than J_i , that is, there exist c higher-priority pending jobs in cluster C_k at time t , which precludes s-oblivious pi-blocking (recall Def. 1). Contradiction. \square

Having bounded maximum pi-blocking during each of the OMIP's "stages" that a job traverses, we obtain the following overall bound on per-resource s-oblivious pi-blocking.

Lemma 7. *J_i incurs at most $b_{i,q} = N_{i,q} \cdot (2m - 1) \cdot L_q^{max}$ s-oblivious pi-blocking due to requests by any task for ℓ_q .*

Proof. By Lemma 1, J_i incurs pi-blocking only while suspended. Consider a single request for ℓ_q . While J_i is suspended waiting to acquire ℓ_q , the amount of pi-blocking incurred is determined by the number of conflicting requests that precede J_i 's request. By Lemma 6, at most m requests must complete before J_i enters FQ_{q,k} or ceases to incur pi-blocking in PQ_{q,k}. By Lemmas 3 and 5, at most $\frac{m}{c} - 1 + \frac{m}{c} \cdot (c - 1) = m - 1$ requests must complete from the time that J_i enters FQ_{q,k} until it holds ℓ_q . Lemma 2 implies that the job holding ℓ_q progresses towards completion of its request whenever J_i incurs s-oblivious pi-blocking. Thus, after J_i has incurred $(m + m - 1) \cdot L_q^{max}$ s-oblivious pi-blocking, at least $2m - 1$ requests must have completed and J_i holds ℓ_q . The stated bound follows from considering the maximum number of requests issued by any J_i for ℓ_q . \square

Note that Lemma 7 implies that the OMIP is indeed independence-preserving (recall Def. 2). Finally, we obtain the following bound on total s-oblivious pi-blocking.

Theorem 2. *Under clustered JLFP scheduling with the OMIP, a job J_i incurs in total at most $b_i = \sum_q b_{i,q} = \sum_q N_{i,q} \cdot (2m - 1) \cdot L_q^{max} = O(m)$ s-oblivious pi-blocking.*

Proof. Follows from Lemma 7 since $b_i = \sum_q b_{i,q}$. \square

Prior work has established $m - 1$ as a simple per-request lower bound under s-oblivious analysis [11]; the OMIP is thus optimal within a factor of $\frac{2m-1}{m-1} = 2 + \frac{1}{m-1} \approx 2$.

5 Implementation and Evaluation

We modified the existing C-EDF plugin in LITMUS^{RT} [1], a real-time extension of the Linux kernel, to incorporate migratory priority inheritance, the OMIP, and the OMLP. The OMIP's multi-level queue structure is easily realized by


combining multiple *wait queues*, which are readily available in Linux. Most complexity stems from the need to coordinate scheduling decisions across cluster boundaries when a job is eligible to be scheduled in multiple clusters due to migratory priority inheritance. In particular, if a processor decides to *not* schedule a lock-holding job J_l because it observes that J_l is already scheduled in another cluster, the processor must be notified if J_l is subsequently preempted. In our prototype, this is realized with a per-semaphore bitmask that tracks on which processors a job holding the semaphore may currently be scheduled, and by sending an *inter-processor interrupt* (IPI) to available processors when a lock-holder preemption occurs. A migration, if possible, is then carried out by the first available processor to reschedule in response to the IPI. A detailed discussion of the implementation is omitted here due to space constraints; however, our changes are publicly available as a patch on the LITMUS^{RT} homepage [1].

5.1 Response-Time Measurements


Based on LITMUS^{RT}, we conducted an experiment to empirically demonstrate that the OMIP is effective at protecting latency-sensitive tasks, that is, to rule out the possibility that independence-preservation is merely an “analytical trick” without practical impact. To this end, we configured a simple synthetic task set on a 2.0 GHz Intel Xeon X7550 system with $m = 8$ cores using the modified LITMUS^{RT} C-EDF plugin with clusters defined by the L1 cache topology, which results in P-EDF scheduling ($c = 1$) on the test platform.

On each core, we launched four tasks with periods $1ms$, $25ms$, $100ms$, and $1000ms$ and execution costs of roughly $0.1ms$, $2ms$, $15ms$, and $600ms$, resp. The latency-sensitive, one-millisecond tasks did not access any shared resources. All other tasks shared a single lock with an associated maximum critical section length of approximately $L^{max} = 1ms$, and each of their jobs acquired the lock once. While the task set is synthetic, the chosen parameters are inspired by the range of periods found in automotive systems [15, 36] and we believe them to be a reasonable approximation of heterogenous timing constraints as they arise in practice.

We ran the task set once using the C-OMLP, once using the OMIP, and once using no locks at all (as a baseline assuming that all tasks are independent) for 30 minutes each. We traced the resulting schedules using LITMUS^{RT}’s `sched_trace` facility and recorded the response times of more than 45,000,000 individual jobs. Fig. 3 shows two histograms of recorded response times under each configuration. Fig. 3(a) depicts response times of the one-millisecond tasks. Due to the short period (and consistent deadline tie-breaking), their jobs always have the highest priority and thus incur no delays at all in the absence of locks (all response times are in the first bin). In contrast, under the C-OMLP, jobs are not protected from pi-blocking due to unrelated critical sections and response times exceeding $8ms$ were observed—priority boosting causes deadline misses in latency-sensitive tasks



(a) Response times of latency-sensitive tasks with period $p_i = 1ms$.



(b) Response times of regular tasks with period $p_i = 100ms$.

Figure 3: Response times measured in LITMUS^{RT} under C-EDF scheduling with the OMIP, the C-OMLP, and without locks.

(not coincidentally, $8ms \approx m \cdot L^{max}$). Not so under the OMIP, where the response-times are identical to the case without locks—task independence was indeed preserved.

However, mutual exclusion invariably causes some delays, and if such delays do not manifest in the highest-priority jobs, then they will be necessarily observable elsewhere. This is apparent in Fig. 3(b), which depicts the response-time distribution of the $100ms$ -tasks. For these tasks with considerable slack, worst-case response-times are noticeably shorter under the C-OMLP than under the OMIP. This emphasizes that there is an obvious tradeoff between not penalizing higher-priority jobs and rapidly completing critical sections. To explore this tradeoff, we conducted schedulability experiments, which allow *analytical* differences w.r.t. response-time *guarantees* to be quantified (as opposed to the differences in *observed* response times discussed so far).

5.2 Schedulability Experiments

In preparation, we derived and implemented fine-grained (*i.e.*, non-asymptotic) pi-blocking analysis of the OMIP suitable for schedulability analysis based on a recently developed analysis technique using linear programming [10]; details are provided in Appendix A.2. To quantify each protocol’s impact on schedulability, we generated task sets using Emberson *et al.*’s method [21] consisting of $n \in \{20, 30, 40\}$ tasks with total utilization $U \in \{0.4m, 0.5m, 0.7m\}$. Of the n tasks, $n^{lat} \in \{0, 1, \dots, 8\}$ were chosen to be latency-sensitive. Latency-sensitive tasks were assigned a period randomly chosen from $[0.5ms, 2.5ms]$, and regular tasks were




Figure 4: The effect of increasing maximum critical section length under P-EDF scheduling with the C-OMLP and the OMIP.

assigned a period randomly chosen from $[10ms, 1000ms]$ (in both cases with a granularity of $500\mu s$).

The n^{lat} latency-sensitive tasks shared three resources, which each job accessed once each with maximum critical section lengths chosen uniformly from $[1\mu s, 15\mu s]$. The remaining $n - n^{lat}$ regular tasks shared 12 different resources, of which each task accessed $N^{max} \in \{1, 2, 3\}$, with a maximum critical section length randomly chosen from $[1, mcsl]$, where $mcsl$ is a variable parameter denoting the maximum critical section length. We varied $mcsl \in [5\mu s, 1000\mu s]$ in steps of $5\mu s$ on processor platforms with $m \in \{2, 4, 8, 16\}$ processors and generated for each combination of n , n^{lat} , U , N^{max} , and $mcsl$ at least 1000 task sets (more than 150,000,000 in total). Each generated task set was schedulability tested under both the C-OMLP and the OMIP assuming P-EDF scheduling (*i.e.*, C-EDF with $c = 1$). The full set of results, 678 schedulability plots when visualized as a function of $mcsl$, is provided in Appendix A.3; three illustrative examples highlighting two main “bottlenecks” that constrain schedulability are shown in Fig. 4.

Fig. 4(a) shows one of many cases where the OMIP is clearly preferable. In this scenario with one latency-sensitive task (*i.e.*, $n^{lat} = 1$, see figure for all parameters), schedulability under the C-OMLP decreases rapidly as critical section lengths start to exceed $\approx 50\mu s$. The cause is that starting with $mcsl = 50\mu s$, latency-sensitive tasks with a period of $p_i = 500\mu s$ are likely rendered unschedulable by delays due to unrelated critical sections in the range of $m \cdot L^{max} \approx 50\mu s = 400\mu s$, similar to the hand-crafted example discussed in Sec. 5.1. In contrast, the OMIP achieves high schedulability even in the presence of much longer critical sections. In this scenario, latency sensitivity is the bottleneck, and the OMIP is preferable for such workloads.

In contrast, Fig. 4(b) shows a scenario with the opposite outcome. In this case, no latency-sensitive tasks are included in the generated task sets ($n^{lat} = 0$) and there is a high degree of contention ($N^{max} = 3$). As a result, independence-preservation is not essential and constant factors in the pi-blocking analysis dominate: while the C-OMLP and the OMIP are both $O(m)$ protocols, the C-OMLP has a per-request bound of $(m - 1) \cdot L^{max}$ [12], whereas the OMIP

has a per-request bound of $(2m - 1) \cdot L^{max}$ (Lemma 7). The C-OMLP thus achieves somewhat higher schedulability in this case, though long critical sections are infeasible under either protocol due to the high level of contention.

Finally, Fig. 4(c) shows a scenario with $n^{lat} = 6$ and $N^{max} = 3$ where both effects combine. Due to the high level of contention, schedulability under the OMIP is limited across the entire range of critical section lengths, whereas under the C-OMLP, schedulability for short critical sections is markedly higher than under the OMIP due to the C-OMLP’s advantage in constant factors, but then decreases rapidly due to the lack of independence preservation under the C-OMLP.

In summary, our results show that independence preservation and the OMIP are clearly beneficial for latency-sensitive workloads. Further, the OMIP is competitive with and often performs similar to the C-OMLP even if independence preservation is not required, but may perform worse than the C-OMLP in the presence of heavy contention.

6 Conclusion

This paper is the first to investigate the challenge of supporting latency-sensitive real-time applications under partitioned and clustered scheduling. We have proposed independence preservation to capture the essential requirement of such workloads. The primary contribution of this paper is the OMIP, the first independence-preserving real-time semaphore protocol for clustered JLFP scheduling. Interestingly, the OMIP employs a novel three-stage queue, which yields asymptotic optimality under s-oblivious analysis.

In summary, latency-sensitive real-time applications have long been well-supported on uniprocessors, but could not be hosted on partitioned multiprocessors with prior semaphore protocols. The OMIP is the first semaphore protocol to overcome this limitation, at the expense of job migrations.

Naturally, job migrations do cause additional overheads, and current industry practice thus regards them critically. However, if job migration is disallowed or fundamentally infeasible (*e.g.*, if the processors in different clusters implement incompatible instruction sets), then Theorem 1 shows that there is little that can be done besides disallowing cross-cluster resource sharing altogether. In contrast, in shared-


memory systems in which global scheduling is technically possible, but a non-global scheduling approach is preferred for other reasons (*e.g.*, to improve cache affinity or to reuse existing uniprocessor algorithms), the OMIP is a viable compromise: first, migratory priority inheritance is triggered only if the lock-holding job is preempted, which means that its cache contents are likely evicted anyway, and second, it is cheaper to migrate critical sections than it is to migrate jobs since the cache footprint of a critical section is typically much smaller than that of an entire job. Concerning the number of times that a critical section may be preempted, migratory priority inheritance is no worse in this regard than the widely deployed PIP. Nonetheless, it will be interesting to study the effects of the OMIP on overheads in more detail.

In future algorithmic work, we seek to lift the assumption that jobs lock only one resource at a time. Ward and Anderson [37] recently showed how to support nested critical sections without loss of asymptotic optimality. However, under s-oblivious analysis, their approach permits at most m concurrent (outermost) critical sections by means of a global token lock, which can serialize otherwise independent jobs. Ward and Anderson’s technique [37] thus preserves asymptotic optimality when applied to the OMIP, but it is not independence-preserving. However, it may be possible to combine Ward and Anderson’s approach with the FMLP’s notion of resource groups [7] to obtain independence preservation and asymptotic optimality despite lock nesting.


Finally, we further seek to explore the design space of independence-preserving locking protocols for s-aware analysis, and are currently investigating synchronization mechanisms that provide “freedom from interference” (*i.e.*, temporal isolation) among mutually untrusting tasks, despite resource sharing in the presence of budget overruns, for which the OMIP is an essential building block [9].

References

- [1] The LITMUS^{RT} project. <http://www.litmus-rt.org>.
- [2] J. Anderson, R. Jain, and K. Jeffay. Efficient object sharing in quantum-based real-time systems. In *Proc. of the 19th Real-Time Systems Symposium*, 1998.
- [3] T. Baker. Stack-based scheduling for realtime processes. *Real-Time Systems*, 3(1):67–99, 1991.
- [4] T. Baker and S. Baruah. Schedulability analysis of multiprocessor sporadic task systems. In *Handbook of Real-Time and Embedded Systems*. Chapman Hall/CRC, 2007.
- [5] S. Baruah. Techniques for multiprocessor global schedulability analysis. In *Proc. of the 28th Real-Time Systems Symposium*, 2007.
- [6] M. Bertogna and M. Cirinei. Response-time analysis for globally scheduled symmetric multiprocessor platforms. In *Proc. of the 28th Real-Time Systems Symposium*, 2007.
- [7] A. Block, H. Leontyev, B. Brandenburg, and J. Anderson. A flexible real-time locking protocol for multiprocessors. In *Proc. of the 13th Conference on Embedded and Real-Time Computing Systems and Applications*, 2007.
- [8] B. Brandenburg. *Scheduling and Locking in Multiprocessor Real-Time Operating Systems*. PhD thesis, The University of North Carolina at Chapel Hill, 2011.
- [9] B. Brandenburg. Virtually exclusive resources. Technical Report MPI-SWS-2012-005, MPI-SWS, May 2012.
- [10] B. Brandenburg. Improved analysis and evaluation of real-time semaphore protocols for P-FP scheduling. In *Proc. of the 19th Real-Time and Embedded Technology and Applications Symposium*, 2013.
- [11] B. Brandenburg and J. Anderson. Optimality results for multiprocessor real-time locking. In *Proc. of the 31st Real-Time Systems Symposium*, 2010.
- [12] B. Brandenburg and J. Anderson. The OMLP family of optimal multiprocessor real-time locking protocols. *Design Automation for Embedded Systems*, online first, July 2012.
- [13] B. Brandenburg and A. Bastoni. The case for migratory priority inheritance in Linux: Bounded priority inversions on multiprocessors. In *Proc. of the 14th Real-Time Linux Workshop*, 2012.
- [14] B. Brandenburg, J. Calandrino, A. Block, H. Leontyev, and J. Anderson. Synchronization on real-time multiprocessors: To block or not to block, to suspend or spin? In *Proc. of the 14th Real-Time and Embedded Technology and Apps. Symposium*, 2008.
- [15] D. Buttle. Real-time in the prime-time. Keynote presentation at the 24th Euromicro Conference on Real-Time Systems, July 2012.
- [16] J. Calandrino, J. Anderson, and D. Baumberger. A hybrid real-time scheduling approach for large-scale multicore platforms. In *Proc. of the 19th Euromicro Conference on Real-Time Systems*, 2007.
- [17] Y. Chang, R. Davis, and A. Wellings. Reducing queue lock pessimism in multiprocessor schedulability analysis. In *Proc. of the 18th Conference on Real-Time and Network Systems*, 2010.
- [18] A. Chou, J. Yang, B. Chelf, S. Hallem, and D. Engler. An empirical study of operating systems errors. In *Proc. of the 18th Symposium on Operating Systems Principles*, 2001.
- [19] T. Craig. Queuing spin lock algorithms to support timing predictability. In *Proc. of the 14th IEEE Real-Time Systems Symposium*, 1993.
- [20] A. Easwaran and B. Andersson. Resource sharing in global fixed-priority preemptive multiprocessor scheduling. In *Proc. of the 30th Real-Time Systems Symposium*, 2009.
- [21] P. Emberson, R. Stafford, and R. Davis. Techniques for the synthesis of multiprocessor tasksets. In *1st Workshop on Analysis Tools and Methodologies for Embedded and Real-time Systems*, 2010.
- [22] Express Logic, Inc. ThreadX real-time operating system. <http://rtos.com/products/threadx/>, accessed 1/26/2013.
- [23] D. Faggioli, G. Lipari, and T. Cucinotta. The multiprocessor bandwidth inheritance protocol. In *Proc. of the 22nd Euromicro Conference on Real-Time Systems*, 2010.
- [24] P. Gai, M. di Natale, G. Lipari, A. Ferrari, C. Gabellini, and P. Marceca. A comparison of MPCP and MSRP when sharing resources in the Janus multiple processor on a chip platform. In *Proc. of the 9th Real-Time And Embedded Technology Application Symposium*, 2003.
- [25] M. Hohmuth and M. Peter. Helping in a multiprocessor environment. In *Proceeding of the Second Workshop on Common Microkernel System Platforms*, 2001.
- [26] G. Macariu and V. Cretu. Limited blocking resource sharing for global multiprocessor scheduling. In *Proc. of the 23rd Euromicro Conference on Real-Time Systems*, 2011.
- [27] F. Nemat, M. Behnam, and T. Nolte. Independently-developed real-time systems on multi-cores with shared resources. In *Proc. of the 23rd Euromicro Conference on Real-Time Systems*, 2011.
- [28] Open Source Automation Development Lab eG. Continuous worst-case latency monitoring. <https://www.osadl.org/Continuous-latency-monitoring.ga-farm-monitoring.0.html>, accessed 1/26/2013.
- [29] R. Rajkumar. Real-time synchronization protocols for shared memory multiprocessors. In *Proc. of the 10th Conference on Dist Computing Systems*, 1990.
- [30] R. Rajkumar. *Synchronization In Real-Time Systems—A Priority Inheritance Approach*. Kluwer Academic Publishers, 1991.
- [31] R. Rajkumar, L. Sha, and J. Lehoczky. Real-time synchronization protocols for multiprocessors. In *Proc. of the 9th Real-Time Systems Symposium*, 1988.
- [32] F. Ridouard, P. Richard, and F. Cottet. Negative results for scheduling independent hard real-time tasks with self-suspensions. In *Proc. of 25th Real-Time Systems Symposium*, 2004.
- [33] L. Sha, R. Rajkumar, and J. Lehoczky. Priority inheritance protocols: an approach to real-time synchronization. *IEEE Trans. Comput.*, 39(9):1175–1185, 1990.
- [34] H. Takada and K. Sakamura. Predictable spin lock algorithms with preemption. In *Proc. of the 11th Workshop on Real-Time Operating Systems and Software*, 1994.
- [35] H. Takada and K. Sakamura. A novel approach to multiprogrammed multiprocessor synchronization for real-time kernels. In *Proc. of the 18th Real-Time Systems Symposium*, 1997.
- [36] K. Tindell and A. Burns. Guaranteeing message latencies on control area network (CAN). In *Proc. of the 1st Intl. CAN Conference*, 1994.
- [37] B. Ward and J. Anderson. Supporting nested locking in multiprocessor real-time systems. In *Proc. of the 24th Euromicro Conference on Real-Time Systems*, 2012.



(a) Response times of tasks with period $p_i = 25ms$.



(b) Response times of tasks with period $p_i = 1000ms$.

Figure 5: Response times measured in LITMUS^{RT} under C-EDF scheduling with the OMIP, the C-OMLP, and without locks.

A Supplemental Material

This appendix provides material that has been omitted from the conference version of this paper due to space constraints. First, Sec. A.1 provides the remaining response-time distributions observed during the experiment discussed in Sec. 5.1. Sec. A.2 provides detailed (*i.e.*, non-asymptotic) blocking analysis of the OMIP, which is used in the schedulability experiments discussed in Sec. 5.2. Finally, Sec. A.3 provides the full set of results for all parameter combinations considered in Sec. 5.2.

A.1 Response-Time Distributions

As discussed in Sec. 5.1, we ran a simple task set consisting of tasks with periods $p_i \in \{1ms, 25ms, 100ms, 1000ms\}$ on an eight-core platform running LITMUS^{RT}. The response-time distributions of tasks with $p_i = 1ms$ and $p_i = 100ms$ were already shown in Fig. 3; the remaining two response-time distributions for the cases $p_i = 25ms$ and $p_i = 1000ms$ are given in Fig. 5. Both distribution exhibit the same trends already discussed in Sec. 5.1: unsurprisingly, the response times are shortest in the absence of blocking, affected by pi-blocking under the C-OMLP, and somewhat larger under the OMIP than under the OMLP.

A.2 Finely-Grained Bounds on PI-Blocking

In this section, we provide a means for deriving non-asymptotic bounds on the maximum s-oblivious pi-blocking

incurred by any job of a given task T_i under the OMLP. The derived bounds are finely-grained in the sense that constant factors such as the number of times that each resource is accessed are fully reflected. As a result, the obtained upper bounds on maximum s-oblivious pi-blocking are considerably less pessimistic than the asymptotic analysis presented in Sec. 4.3 and thus better suited to schedulability analysis.

Definitions. Some additional notation is required for the following discussion. Let J_i denote an arbitrary job of the task under analysis. For each resource ℓ_q and other task T_x (*i.e.*, $T_i \neq T_x$), we let $N_{x,q}^i$ denote the maximum number of requests issued by jobs of T_x for ℓ_q while J_i is pending. This bound can be derived from the parameter $N_{x,q}$ and the maximum number of jobs of T_x that can overlap with any J_i , which depends on the *maximum response times* of T_i and T_x (*i.e.*, the maximum time that any job remains pending). Let r_i (r_x , resp.) denote the maximum response-time of T_i (T_x , resp.); then $N_{x,q}^i$ is defined as follows [8, 10]:

$$N_{x,q}^i = N_{x,q} \cdot \left\lceil \frac{r_i + r_x}{p_x} \right\rceil.$$

Additionally, we let $A_{k,q}$ denote the number of tasks assigned to cluster C_k that access ℓ_q ; formally

$$A_{k,q} = |\{T_j \mid C(T_j) = C_k \wedge N_{j,q} > 0\}|.$$

Finally, we assume a slightly modified version of the OMIP that has improved constant factors in certain special cases. As shown in Sec. 4.3, key to the OMIP's asymptotic optimality is that a job will incur s-oblivious pi-blocking for at most the duration that it takes $2c - 1$ jobs to be dequeued from $FQ_{q,k}$ *regardless of* $A_{k,q}$, which is ensured by the hybrid queue employed in each cluster (*i.e.*, the combination of $PQ_{q,k}$ and $FQ_{q,k}$). If, however, $A_{k,q} \leq 2c$ in some cluster C_k , then a simple FIFO queue is sufficient, which simplifies the analysis somewhat. We therefore assume the following updated request rule (and that each $FQ_{q,k}$ can hold $2c$ jobs).

R1' When J_i issues a request for ℓ_q and $FQ_{q,k}$ is empty, then J_i is enqueued in both GQ_q and $FQ_{q,k}$. Otherwise, if there are fewer than c jobs queued in $FQ_{q,k}$ or if $A_{k,q} \leq 2c$, then J_i is enqueued only in $FQ_{q,k}$. Finally, if there are already c jobs queued in $FQ_{q,k}$ and $A_{k,q} > 2c$, then J_i is enqueued in $PQ_{q,k}$.

Note that the asymptotic complexity is not affected by Rule R1' (*i.e.*, this constant-factor optimization is transparent to the asymptotic analysis presented in Sec. 4.3).

LP-based analysis. In the following, to bound maximum pi-blocking b_i for a task T_i , we derive a linear program (LP) that, when maximized by an LP solver, yields a safe upper bound on b_i . This technique was first proposed and explained in detail in [10]. In short, it works as follows. For each T_x , ℓ_q , and each of the at most $N_{x,q}^i$ concurrent

requests of T_x for ℓ_q , we define a variable $X_{x,q,v}$ with domain $[0, 1]$, where $v \in \{1, \dots, N_{x,q}^i\}$. Each such variable, also called *blocking fractions*, corresponds the fraction of the respective maximum critical section length that contributes to J_i 's pi-blocking. The bound on T_i 's maximum s-oblivious pi-blocking b_i can then be expressed as follows.

$$b_i = \sum_{T_x \neq T_i} \sum_{q=1}^r \sum_{v=1}^{N_{x,q}^i} L_{x,q} \cdot X_{x,q,v} \quad (1)$$

Eq. (1) is the objective function of the generated LP for task T_i , that is, the objective of the LP solver is to find a combination of blocking critical sections that maximize b_i . To avoid excessive pessimism, additional linear constraints that encode properties of the OMIP are imposed on blocking fractions. The following constraints prevent the LP solver from considering clearly impossible scenarios.

As shown in Sec. 4.3, the OMIP is independence-preserving and ensures that J_i incurs s-oblivious pi-blocking for the cumulative duration of at most $2m - 1$ critical section lengths due to each request that it issues. This basic bound on per-request pi-blocking implies the following constraint.

Constraint 1. *In any JLFP schedule of τ under the OMIP:*

$$\forall \ell_q : \sum_{T_x \neq T_i} \sum_{v=1}^{N_{x,q}^i} X_{x,q,v} \leq N_{i,q} \cdot (2m - 1).$$

Next, we consider pi-blocking due to requests issued by tasks local to T_i . In the case that $A_{i,q} \leq 2c$, then the use of FIFO queues only implies that each other task can block J_i with at most one request. In the case that $A_{i,q} > 2c$, no strict FIFO order exists; however, it can be shown that each task contributes pi-blocking with at most two critical sections [8]. Let $\tau^l \triangleq \{T_x \mid C(T_x) = C(T_i) \wedge T_x \neq T_i\}$ denote the set of local tasks w.r.t. T_i , and let $C_k = C(T_i)$. The following constraint results.

Constraint 2. *In any JLFP schedule of τ under the OMIP:*

$$\forall T_x \in \tau^l : \forall \ell_q : \sum_{v=1}^{N_{x,q}^i} X_{x,q,v} \leq \begin{cases} N_{i,q} & \text{if } A_{k,q} \leq 2c, \\ 2N_{i,q} & \text{otherwise.} \end{cases}$$

Finally, we bound the maximum contribution of remote tasks in each remote cluster. Remote tasks can only block J_i by preceding another job from a local task in the global queue GQ_q . In the case of $A_{k,q} \leq 2c$, the exclusive use of FIFO queues implies an upper bound of $A_{k,q}$ blocking critical sections from each remote cluster per each of J_i 's requests (since at most $A_{k,q}$ jobs from $C(T_i)$ must traverse GQ_q before J_i holds ℓ_q). In the case of $A_{k,q} > 2c$, a somewhat coarser bound applies: while J_i traverses $FQ_{q,k}$, tasks in each remote cluster can block J_i with at most c critical sections, and while J_i waits in $PQ_{q,k}$, J_i is blocked by

at most m critical sections (recall Lemma 6). These considerations lead to the following constraint. Let $\tau(C_y) \triangleq \{T_x \mid C(T_x) = C_y\}$ denote the set of tasks assigned to remote cluster C_y .

Constraint 3. *In any JLFP schedule of τ under the OMIP:*

$$\forall \ell_q : \forall C_y \neq C(T_i) :$$

$$\sum_{T_x \in \tau(C_y)} \sum_{v=1}^{N_{x,q}^i} X_{x,q,v} \leq \begin{cases} N_{i,q} \cdot A_{k,q} & \text{if } A_{k,q} \leq 2c, \\ N_{i,q} \cdot (c + m) & \text{otherwise.} \end{cases}$$

The above analysis was implemented in the open-source *Schedulability test Collection And Toolkit*⁵ (SchedCAT), which was used to carry out the experiments reported on in Sec. 5.2. SchedCAT was configured to use the *GNU Linear Programming Toolkit*⁶ (GLPK) to solve each generated LP. When bounding each $N_{x,q}^i$, we approximated r_i and r_x with p_i and p_x , resp., since the (implicit) relative deadline of each task represents an upper bound on the maximum response time if the task set is found to be schedulable.

A.3 Schedulability Results

The schedulability graphs for all 678 configurations considered in the schedulability experiments described in Sec. 5.2 are shown in Figs. 6–683.

Organization. The presented graphs are organized by parameters as outlined below; a classification of the graphs by major trends is given thereafter.

- Figs. 6–86 show results for $m = 2$.
 - Figs. 6–32 show results for $n = 20$.
 - * Figs. 6–14 show results for $n^{lat} = 0$.
 - * Figs. 15–23 show results for $n^{lat} = 1$.
 - * Figs. 24–32 show results for $n^{lat} = 2$.
 - Figs. 33–59 show results for $n = 30$.
 - * Figs. 33–41 show results for $n^{lat} = 0$.
 - * Figs. 42–50 show results for $n^{lat} = 1$.
 - * Figs. 51–59 show results for $n^{lat} = 2$.
 - Figs. 60–86 show results for $n = 40$.
 - * Figs. 60–68 show results for $n^{lat} = 0$.
 - * Figs. 69–77 show results for $n^{lat} = 1$.
 - * Figs. 78–86 show results for $n^{lat} = 2$.
- Figs. 87–221 show results for $m = 4$.
 - Figs. 87–131 show results for $n = 20$.
 - * Figs. 87–95 show results for $n^{lat} = 0$.
 - * Figs. 96–104 show results for $n^{lat} = 1$.

⁵<http://www.mpi-sws.org/~bbb/projects/schedcat>

⁶<http://www.gnu.org/software/glpk>

- * Figs. 105–113 show results for $n^{lat} = 2$.
- * Figs. 114–122 show results for $n^{lat} = 3$.
- * Figs. 123–131 show results for $n^{lat} = 4$.
- Figs. 132–176 show results for $n = 30$.
 - * Figs. 132–140 show results for $n^{lat} = 0$.
 - * Figs. 141–149 show results for $n^{lat} = 1$.
 - * Figs. 150–158 show results for $n^{lat} = 2$.
 - * Figs. 159–167 show results for $n^{lat} = 3$.
 - * Figs. 168–176 show results for $n^{lat} = 4$.
- Figs. 177–221 show results for $n = 40$.
 - * Figs. 177–185 show results for $n^{lat} = 0$.
 - * Figs. 186–194 show results for $n^{lat} = 1$.
 - * Figs. 195–203 show results for $n^{lat} = 2$.
 - * Figs. 204–212 show results for $n^{lat} = 3$.
 - * Figs. 213–221 show results for $n^{lat} = 4$.
- Figs. 222–464 show results for $m = 8$.
 - Figs. 222–302 show results for $n = 20$.
 - * Figs. 222–230 show results for $n^{lat} = 0$.
 - * Figs. 231–239 show results for $n^{lat} = 1$.
 - * Figs. 240–248 show results for $n^{lat} = 2$.
 - * Figs. 249–257 show results for $n^{lat} = 3$.
 - * Figs. 258–266 show results for $n^{lat} = 4$.
 - * Figs. 267–275 show results for $n^{lat} = 5$.
 - * Figs. 276–284 show results for $n^{lat} = 6$.
 - * Figs. 285–293 show results for $n^{lat} = 7$.
 - * Figs. 294–302 show results for $n^{lat} = 8$.
 - Figs. 303–383 show results for $n = 30$.
 - * Figs. 303–311 show results for $n^{lat} = 0$.
 - * Figs. 312–320 show results for $n^{lat} = 1$.
 - * Figs. 321–329 show results for $n^{lat} = 2$.
 - * Figs. 330–338 show results for $n^{lat} = 3$.
 - * Figs. 339–347 show results for $n^{lat} = 4$.
 - * Figs. 348–356 show results for $n^{lat} = 5$.
 - * Figs. 357–365 show results for $n^{lat} = 6$.
 - * Figs. 366–374 show results for $n^{lat} = 7$.
 - * Figs. 375–383 show results for $n^{lat} = 8$.
 - Figs. 384–464 show results for $n = 40$.
 - * Figs. 384–392 show results for $n^{lat} = 0$.
 - * Figs. 393–401 show results for $n^{lat} = 1$.
 - * Figs. 402–410 show results for $n^{lat} = 2$.
 - * Figs. 411–419 show results for $n^{lat} = 3$.
 - * Figs. 420–428 show results for $n^{lat} = 4$.
 - * Figs. 429–437 show results for $n^{lat} = 5$.
 - * Figs. 438–446 show results for $n^{lat} = 6$.
 - * Figs. 447–455 show results for $n^{lat} = 7$.
- * Figs. 456–464 show results for $n^{lat} = 8$.
- Figs. 465–683 show results for $m = 16$.
 - Figs. 465–530 show results for $n = 20$.
 - * Figs. 465–473 show results for $n^{lat} = 0$.
 - * Figs. 474–482 show results for $n^{lat} = 1$.
 - * Figs. 483–491 show results for $n^{lat} = 2$.
 - * Figs. 492–500 show results for $n^{lat} = 3$.
 - * Figs. 501–509 show results for $n^{lat} = 4$.
 - * Figs. 510–518 show results for $n^{lat} = 5$.
 - * Figs. 519–524 show results for $n^{lat} = 6$.
 - * Figs. 525–530 show results for $n^{lat} = 7$.
 - Figs. 531–602 show results for $n = 30$.
 - * Figs. 531–539 show results for $n^{lat} = 0$.
 - * Figs. 540–548 show results for $n^{lat} = 1$.
 - * Figs. 549–557 show results for $n^{lat} = 2$.
 - * Figs. 558–566 show results for $n^{lat} = 3$.
 - * Figs. 567–575 show results for $n^{lat} = 4$.
 - * Figs. 576–584 show results for $n^{lat} = 5$.
 - * Figs. 585–593 show results for $n^{lat} = 6$.
 - * Figs. 594–602 show results for $n^{lat} = 7$.
 - Figs. 603–683 show results for $n = 40$.
 - * Figs. 603–611 show results for $n^{lat} = 0$.
 - * Figs. 612–620 show results for $n^{lat} = 1$.
 - * Figs. 621–629 show results for $n^{lat} = 2$.
 - * Figs. 630–638 show results for $n^{lat} = 3$.
 - * Figs. 639–647 show results for $n^{lat} = 4$.
 - * Figs. 648–656 show results for $n^{lat} = 5$.
 - * Figs. 657–665 show results for $n^{lat} = 6$.
 - * Figs. 666–674 show results for $n^{lat} = 7$.
 - * Figs. 675–683 show results for $n^{lat} = 16$.

Summary of trends. We classified the results into three major categories: scenarios in which the OMLP provides clearly higher schedulability; scenarios in which the OMIP provides clearly higher schedulability; and scenarios in which neither protocol dominates the other for all critical section lengths.

- Schedulability under the OMLP dominates schedulability under the OMIP in the following 71 scenarios: Figs. 7–11, 13, 14, 34, 35, 37, 38, 40, 41, 44, 47, 61, 62, 64, 65, 67, 68, 71, 74, 77, 89, 92, 95, 98, 133, 134, 136, 137, 139, 140, 142, 143, 178, 179, 181, 182, 184, 185, 187, 188, 224, 227, 230, 233, 294, 295–302, 305, 308, 311, 314, 386, 389, 392, 394, 395, 605, 608, 611, 614, 617.
- Schedulability under the OMIP dominates schedulability under the OMLP in the following 426 scenarios: Figs. 6, 12, 15, 16–33, 36, 39, 42, 43, 45, 46, 48, 49–60,

63, 66, 69, 72, 75, 76, 78, 79–88, 90, 91, 93, 94, 96, 99, 100–125, 132, 135, 138, 141, 144, 145–170, 177, 180, 183, 186, 189, 190, 192, 193–196, 198, 199, 201, 202–215, 223, 225, 226, 228, 229, 232, 234, 235–263, 303, 304, 306, 307, 309, 310, 312, 313, 315, 316–350, 364, 365, 384, 387, 390, 393, 396, 397–437, 444, 445–446, 480, 481–500, 532, 534, 535, 537, 538, 541, 543, 544, 546, 547–590, 594, 595–599, 603, 604, 606, 607, 609, 610, 612, 613, 615, 616, 618, 619–674.

- In the following 181 scenarios neither the OMLP nor the OMIP dominate the other: Figs. 70, 73, 97, 126, 127–131, 171, 172–176, 191, 197, 200, 216, 217–222, 231, 264, 265–293, 351, 352–363, 366, 367–383, 385, 388, 391, 438, 439–443, 447, 448–479, 501, 502–531, 533, 536, 539, 540, 542, 545, 591, 592–593, 600, 601–602, 675, 676–683.

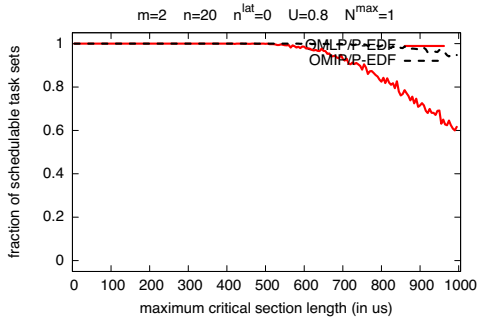


Figure 6: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 0$, $U = 0.80$, and $N^{max} = 1$.




Figure 10: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 0$, $U = 1.00$, and $N^{max} = 2$.




Figure 7: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 0$, $U = 0.80$, and $N^{max} = 2$.




Figure 11: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 0$, $U = 1.00$, and $N^{max} = 3$.




Figure 8: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 0$, $U = 0.80$, and $N^{max} = 3$.




Figure 12: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 0$, $U = 1.40$, and $N^{max} = 1$.




Figure 9: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 0$, $U = 1.00$, and $N^{max} = 1$.




Figure 13: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 0$, $U = 1.40$, and $N^{max} = 2$.




Figure 14: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 0$, $U = 1.40$, and $N^{max} = 3$.




Figure 18: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 1$, $U = 1.00$, and $N^{max} = 1$.

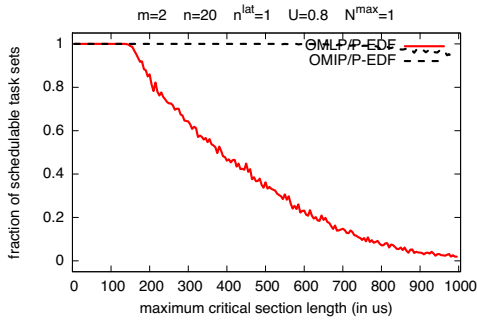


Figure 15: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 1$, $U = 0.80$, and $N^{max} = 1$.




Figure 19: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 1$, $U = 1.00$, and $N^{max} = 2$.




Figure 16: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 1$, $U = 0.80$, and $N^{max} = 2$.

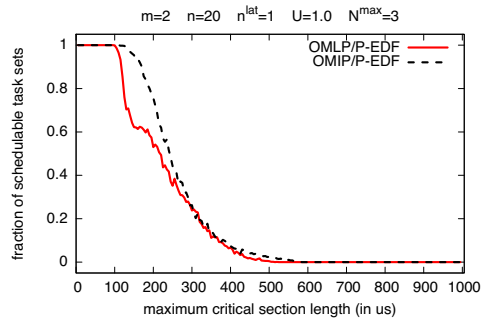


Figure 20: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 1$, $U = 1.00$, and $N^{max} = 3$.




Figure 17: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 1$, $U = 0.80$, and $N^{max} = 3$.




Figure 21: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 1$, $U = 1.40$, and $N^{max} = 1$.




Figure 22: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 1$, $U = 1.40$, and $N^{max} = 2$.

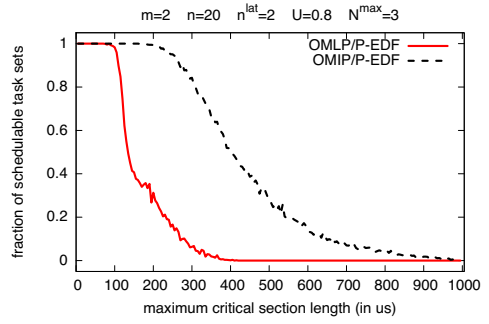


Figure 26: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 2$, $U = 0.80$, and $N^{max} = 3$.




Figure 23: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 1$, $U = 1.40$, and $N^{max} = 3$.




Figure 27: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 2$, $U = 1.00$, and $N^{max} = 1$.




Figure 24: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 2$, $U = 0.80$, and $N^{max} = 1$.




Figure 28: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 2$, $U = 1.00$, and $N^{max} = 2$.




Figure 25: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 2$, $U = 0.80$, and $N^{max} = 2$.




Figure 29: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 2$, $U = 1.00$, and $N^{max} = 3$.




Figure 30: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 2$, $U = 1.40$, and $N^{max} = 1$.




Figure 34: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 0$, $U = 0.80$, and $N^{max} = 2$.




Figure 31: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 2$, $U = 1.40$, and $N^{max} = 2$.




Figure 35: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 0$, $U = 0.80$, and $N^{max} = 3$.




Figure 32: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 20$, $n^{lat} = 2$, $U = 1.40$, and $N^{max} = 3$.




Figure 36: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 0$, $U = 1.00$, and $N^{max} = 1$.




Figure 33: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 0$, $U = 0.80$, and $N^{max} = 1$.




Figure 37: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 0$, $U = 1.00$, and $N^{max} = 2$.




Figure 38: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 0$, $U = 1.00$, and $N^{max} = 3$.




Figure 42: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 1$, $U = 0.80$, and $N^{max} = 1$.




Figure 39: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 0$, $U = 1.40$, and $N^{max} = 1$.




Figure 43: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 1$, $U = 0.80$, and $N^{max} = 2$.




Figure 40: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 0$, $U = 1.40$, and $N^{max} = 2$.




Figure 44: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 1$, $U = 0.80$, and $N^{max} = 3$.




Figure 41: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 0$, $U = 1.40$, and $N^{max} = 3$.




Figure 45: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 1$, $U = 1.00$, and $N^{max} = 1$.




Figure 46: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 1$, $U = 1.00$, and $N^{max} = 2$.




Figure 50: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 1$, $U = 1.40$, and $N^{max} = 3$.




Figure 47: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 1$, $U = 1.00$, and $N^{max} = 3$.




Figure 51: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 2$, $U = 0.80$, and $N^{max} = 1$.




Figure 48: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 1$, $U = 1.40$, and $N^{max} = 1$.




Figure 52: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 2$, $U = 0.80$, and $N^{max} = 2$.




Figure 49: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 1$, $U = 1.40$, and $N^{max} = 2$.




Figure 53: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 2$, $U = 0.80$, and $N^{max} = 3$.




Figure 54: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 2$, $U = 1.00$, and $N^{max} = 1$.




Figure 58: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 2$, $U = 1.40$, and $N^{max} = 2$.




Figure 55: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 2$, $U = 1.00$, and $N^{max} = 2$.




Figure 59: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 2$, $U = 1.40$, and $N^{max} = 3$.




Figure 56: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 2$, $U = 1.00$, and $N^{max} = 3$.




Figure 60: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 0$, $U = 0.80$, and $N^{max} = 1$.




Figure 57: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 30$, $n^{lat} = 2$, $U = 1.40$, and $N^{max} = 1$.




Figure 61: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 0$, $U = 0.80$, and $N^{max} = 2$.




Figure 62: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 0$, $U = 0.80$, and $N^{max} = 3$.

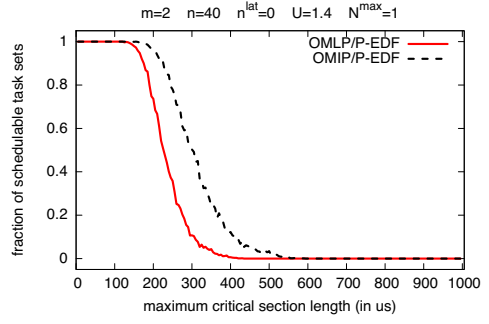


Figure 66: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 0$, $U = 1.40$, and $N^{max} = 1$.




Figure 63: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 0$, $U = 1.00$, and $N^{max} = 1$.




Figure 67: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 0$, $U = 1.40$, and $N^{max} = 2$.




Figure 64: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 0$, $U = 1.00$, and $N^{max} = 2$.




Figure 68: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 0$, $U = 1.40$, and $N^{max} = 3$.




Figure 65: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 0$, $U = 1.00$, and $N^{max} = 3$.




Figure 69: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 1$, $U = 0.80$, and $N^{max} = 1$.




Figure 70: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 1$, $U = 0.80$, and $N^{max} = 2$.




Figure 74: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 1$, $U = 1.00$, and $N^{max} = 3$.




Figure 71: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 1$, $U = 0.80$, and $N^{max} = 3$.




Figure 75: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 1$, $U = 1.40$, and $N^{max} = 1$.




Figure 72: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 1$, $U = 1.00$, and $N^{max} = 1$.




Figure 76: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 1$, $U = 1.40$, and $N^{max} = 2$.




Figure 73: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 1$, $U = 1.00$, and $N^{max} = 2$.




Figure 77: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 1$, $U = 1.40$, and $N^{max} = 3$.




Figure 78: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 2$, $U = 0.80$, and $N^{max} = 1$.




Figure 82: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 2$, $U = 1.00$, and $N^{max} = 2$.




Figure 79: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 2$, $U = 0.80$, and $N^{max} = 2$.




Figure 83: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 2$, $U = 1.00$, and $N^{max} = 3$.




Figure 80: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 2$, $U = 0.80$, and $N^{max} = 3$.




Figure 84: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 2$, $U = 1.40$, and $N^{max} = 1$.




Figure 81: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 2$, $U = 1.00$, and $N^{max} = 1$.




Figure 85: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 2$, $U = 1.40$, and $N^{max} = 2$.




Figure 86: Schedulability under the OMLP and the OMIP for $m = 2$, $n = 40$, $n^{lat} = 2$, $U = 1.40$, and $N^{max} = 3$.




Figure 90: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 0$, $U = 2.00$, and $N^{max} = 1$.




Figure 87: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 0$, $U = 1.60$, and $N^{max} = 1$.




Figure 91: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 0$, $U = 2.00$, and $N^{max} = 2$.




Figure 88: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 0$, $U = 1.60$, and $N^{max} = 2$.




Figure 92: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 0$, $U = 2.00$, and $N^{max} = 3$.




Figure 89: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 0$, $U = 1.60$, and $N^{max} = 3$.




Figure 93: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 0$, $U = 2.80$, and $N^{max} = 1$.




Figure 94: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 0$, $U = 2.80$, and $N^{max} = 2$.




Figure 98: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 1$, $U = 1.60$, and $N^{max} = 3$.




Figure 95: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 0$, $U = 2.80$, and $N^{max} = 3$.




Figure 99: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 1$, $U = 2.00$, and $N^{max} = 1$.




Figure 96: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 1$, $U = 1.60$, and $N^{max} = 1$.




Figure 100: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 1$, $U = 2.00$, and $N^{max} = 2$.




Figure 97: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 1$, $U = 1.60$, and $N^{max} = 2$.




Figure 101: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 1$, $U = 2.00$, and $N^{max} = 3$.




Figure 102: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 1$, $U = 2.80$, and $N^{max} = 1$.




Figure 106: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 2$, $U = 1.60$, and $N^{max} = 2$.




Figure 103: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 1$, $U = 2.80$, and $N^{max} = 2$.




Figure 107: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 2$, $U = 1.60$, and $N^{max} = 3$.




Figure 104: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 1$, $U = 2.80$, and $N^{max} = 3$.




Figure 108: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 2$, $U = 2.00$, and $N^{max} = 1$.




Figure 105: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 2$, $U = 1.60$, and $N^{max} = 1$.




Figure 109: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 2$, $U = 2.00$, and $N^{max} = 2$.




Figure 110: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 2$, $U = 2.00$, and $N^{max} = 3$.




Figure 114: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 3$, $U = 1.60$, and $N^{max} = 1$.




Figure 111: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 2$, $U = 2.80$, and $N^{max} = 1$.




Figure 115: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 3$, $U = 1.60$, and $N^{max} = 2$.




Figure 112: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 2$, $U = 2.80$, and $N^{max} = 2$.




Figure 116: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 3$, $U = 1.60$, and $N^{max} = 3$.




Figure 113: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 2$, $U = 2.80$, and $N^{max} = 3$.




Figure 117: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 3$, $U = 2.00$, and $N^{max} = 1$.




Figure 118: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 3$, $U = 2.00$, and $N^{max} = 2$.




Figure 122: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 3$, $U = 2.80$, and $N^{max} = 3$.




Figure 119: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 3$, $U = 2.00$, and $N^{max} = 3$.




Figure 123: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 4$, $U = 1.60$, and $N^{max} = 1$.




Figure 120: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 3$, $U = 2.80$, and $N^{max} = 1$.




Figure 124: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 4$, $U = 1.60$, and $N^{max} = 2$.




Figure 121: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 3$, $U = 2.80$, and $N^{max} = 2$.




Figure 125: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 4$, $U = 1.60$, and $N^{max} = 3$.




Figure 126: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 4$, $U = 2.00$, and $N^{max} = 1$.




Figure 130: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 4$, $U = 2.80$, and $N^{max} = 2$.




Figure 127: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 4$, $U = 2.00$, and $N^{max} = 2$.




Figure 131: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 4$, $U = 2.80$, and $N^{max} = 3$.




Figure 128: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 4$, $U = 2.00$, and $N^{max} = 3$.




Figure 132: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 0$, $U = 1.60$, and $N^{max} = 1$.




Figure 129: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 20$, $n^{lat} = 4$, $U = 2.80$, and $N^{max} = 1$.




Figure 133: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 0$, $U = 1.60$, and $N^{max} = 2$.




Figure 134: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 0$, $U = 1.60$, and $N^{max} = 3$.




Figure 138: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 0$, $U = 2.80$, and $N^{max} = 1$.




Figure 135: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 0$, $U = 2.00$, and $N^{max} = 1$.




Figure 139: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 0$, $U = 2.80$, and $N^{max} = 2$.




Figure 136: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 0$, $U = 2.00$, and $N^{max} = 2$.




Figure 140: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 0$, $U = 2.80$, and $N^{max} = 3$.




Figure 137: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 0$, $U = 2.00$, and $N^{max} = 3$.




Figure 141: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 1$, $U = 1.60$, and $N^{max} = 1$.




Figure 142: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 1$, $U = 1.60$, and $N^{max} = 2$.




Figure 146: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 1$, $U = 2.00$, and $N^{max} = 3$.




Figure 143: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 1$, $U = 1.60$, and $N^{max} = 3$.




Figure 147: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 1$, $U = 2.80$, and $N^{max} = 1$.




Figure 144: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 1$, $U = 2.00$, and $N^{max} = 1$.




Figure 148: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 1$, $U = 2.80$, and $N^{max} = 2$.




Figure 145: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 1$, $U = 2.00$, and $N^{max} = 2$.




Figure 149: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 1$, $U = 2.80$, and $N^{max} = 3$.




Figure 150: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 2$, $U = 1.60$, and $N^{max} = 1$.




Figure 154: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 2$, $U = 2.00$, and $N^{max} = 2$.




Figure 151: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 2$, $U = 1.60$, and $N^{max} = 2$.




Figure 155: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 2$, $U = 2.00$, and $N^{max} = 3$.




Figure 152: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 2$, $U = 1.60$, and $N^{max} = 3$.




Figure 156: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 2$, $U = 2.80$, and $N^{max} = 1$.




Figure 153: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 2$, $U = 2.00$, and $N^{max} = 1$.




Figure 157: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 2$, $U = 2.80$, and $N^{max} = 2$.




Figure 158: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 2$, $U = 2.80$, and $N^{max} = 3$.




Figure 162: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 3$, $U = 2.00$, and $N^{max} = 1$.




Figure 159: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 3$, $U = 1.60$, and $N^{max} = 1$.




Figure 163: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 3$, $U = 2.00$, and $N^{max} = 2$.




Figure 160: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 3$, $U = 1.60$, and $N^{max} = 2$.




Figure 164: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 3$, $U = 2.00$, and $N^{max} = 3$.




Figure 161: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 3$, $U = 1.60$, and $N^{max} = 3$.




Figure 165: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 3$, $U = 2.80$, and $N^{max} = 1$.




Figure 166: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 3$, $U = 2.80$, and $N^{max} = 2$.




Figure 170: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 4$, $U = 1.60$, and $N^{max} = 3$.




Figure 167: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 3$, $U = 2.80$, and $N^{max} = 3$.




Figure 171: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 4$, $U = 2.00$, and $N^{max} = 1$.




Figure 168: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 4$, $U = 1.60$, and $N^{max} = 1$.




Figure 172: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 4$, $U = 2.00$, and $N^{max} = 2$.




Figure 169: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 4$, $U = 1.60$, and $N^{max} = 2$.




Figure 173: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 4$, $U = 2.00$, and $N^{max} = 3$.




Figure 174: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 4$, $U = 2.80$, and $N^{max} = 1$.




Figure 178: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 0$, $U = 1.60$, and $N^{max} = 2$.




Figure 175: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 4$, $U = 2.80$, and $N^{max} = 2$.




Figure 179: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 0$, $U = 1.60$, and $N^{max} = 3$.




Figure 176: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 30$, $n^{lat} = 4$, $U = 2.80$, and $N^{max} = 3$.




Figure 180: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 0$, $U = 2.00$, and $N^{max} = 1$.




Figure 177: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 0$, $U = 1.60$, and $N^{max} = 1$.




Figure 181: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 0$, $U = 2.00$, and $N^{max} = 2$.




Figure 182: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 0$, $U = 2.00$, and $N^{max} = 3$.




Figure 186: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 1$, $U = 1.60$, and $N^{max} = 1$.




Figure 183: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 0$, $U = 2.80$, and $N^{max} = 1$.




Figure 187: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 1$, $U = 1.60$, and $N^{max} = 2$.




Figure 184: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 0$, $U = 2.80$, and $N^{max} = 2$.




Figure 188: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 1$, $U = 1.60$, and $N^{max} = 3$.




Figure 185: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 0$, $U = 2.80$, and $N^{max} = 3$.




Figure 189: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 1$, $U = 2.00$, and $N^{max} = 1$.




Figure 190: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 1$, $U = 2.00$, and $N^{max} = 2$.




Figure 194: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 1$, $U = 2.80$, and $N^{max} = 3$.




Figure 191: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 1$, $U = 2.00$, and $N^{max} = 3$.




Figure 195: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 2$, $U = 1.60$, and $N^{max} = 1$.




Figure 192: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 1$, $U = 2.80$, and $N^{max} = 1$.




Figure 196: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 2$, $U = 1.60$, and $N^{max} = 2$.




Figure 193: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 1$, $U = 2.80$, and $N^{max} = 2$.




Figure 197: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 2$, $U = 1.60$, and $N^{max} = 3$.




Figure 198: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 2$, $U = 2.00$, and $N^{max} = 1$.




Figure 202: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 2$, $U = 2.80$, and $N^{max} = 2$.




Figure 199: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 2$, $U = 2.00$, and $N^{max} = 2$.




Figure 203: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 2$, $U = 2.80$, and $N^{max} = 3$.




Figure 200: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 2$, $U = 2.00$, and $N^{max} = 3$.




Figure 204: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 3$, $U = 1.60$, and $N^{max} = 1$.




Figure 201: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 2$, $U = 2.80$, and $N^{max} = 1$.




Figure 205: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 3$, $U = 1.60$, and $N^{max} = 2$.




Figure 206: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 3$, $U = 1.60$, and $N^{max} = 3$.




Figure 210: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 3$, $U = 2.80$, and $N^{max} = 1$.




Figure 207: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 3$, $U = 2.00$, and $N^{max} = 1$.




Figure 211: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 3$, $U = 2.80$, and $N^{max} = 2$.




Figure 208: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 3$, $U = 2.00$, and $N^{max} = 2$.




Figure 212: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 3$, $U = 2.80$, and $N^{max} = 3$.




Figure 209: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 3$, $U = 2.00$, and $N^{max} = 3$.




Figure 213: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 4$, $U = 1.60$, and $N^{max} = 1$.




Figure 214: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 4$, $U = 1.60$, and $N^{max} = 2$.




Figure 218: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 4$, $U = 2.00$, and $N^{max} = 3$.




Figure 215: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 4$, $U = 1.60$, and $N^{max} = 3$.




Figure 219: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 4$, $U = 2.80$, and $N^{max} = 1$.




Figure 216: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 4$, $U = 2.00$, and $N^{max} = 1$.




Figure 220: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 4$, $U = 2.80$, and $N^{max} = 2$.




Figure 217: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 4$, $U = 2.00$, and $N^{max} = 2$.




Figure 221: Schedulability under the OMLP and the OMIP for $m = 4$, $n = 40$, $n^{lat} = 4$, $U = 2.80$, and $N^{max} = 3$.




Figure 222: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 0$, $U = 3.20$, and $N^{max} = 1$.




Figure 226: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 0$, $U = 4.00$, and $N^{max} = 2$.




Figure 223: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 0$, $U = 3.20$, and $N^{max} = 2$.




Figure 227: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 0$, $U = 4.00$, and $N^{max} = 3$.




Figure 224: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 0$, $U = 3.20$, and $N^{max} = 3$.




Figure 228: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 0$, $U = 5.60$, and $N^{max} = 1$.




Figure 225: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 0$, $U = 4.00$, and $N^{max} = 1$.




Figure 229: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 0$, $U = 5.60$, and $N^{max} = 2$.




Figure 230: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 0$, $U = 5.60$, and $N^{max} = 3$.




Figure 234: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 1$, $U = 4.00$, and $N^{max} = 1$.




Figure 231: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 1$, $U = 3.20$, and $N^{max} = 1$.




Figure 235: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 1$, $U = 4.00$, and $N^{max} = 2$.




Figure 232: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 1$, $U = 3.20$, and $N^{max} = 2$.




Figure 236: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 1$, $U = 4.00$, and $N^{max} = 3$.




Figure 233: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 1$, $U = 3.20$, and $N^{max} = 3$.




Figure 237: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 1$, $U = 5.60$, and $N^{max} = 1$.




Figure 238: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 1$, $U = 5.60$, and $N^{max} = 2$.




Figure 242: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 2$, $U = 3.20$, and $N^{max} = 3$.




Figure 239: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 1$, $U = 5.60$, and $N^{max} = 3$.




Figure 243: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 2$, $U = 4.00$, and $N^{max} = 1$.




Figure 240: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 2$, $U = 3.20$, and $N^{max} = 1$.




Figure 244: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 2$, $U = 4.00$, and $N^{max} = 2$.




Figure 241: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 2$, $U = 3.20$, and $N^{max} = 2$.




Figure 245: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 2$, $U = 4.00$, and $N^{max} = 3$.




Figure 246: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 2$, $U = 5.60$, and $N^{max} = 1$.




Figure 250: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 3$, $U = 3.20$, and $N^{max} = 2$.




Figure 247: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 2$, $U = 5.60$, and $N^{max} = 2$.




Figure 251: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 3$, $U = 3.20$, and $N^{max} = 3$.




Figure 248: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 2$, $U = 5.60$, and $N^{max} = 3$.




Figure 252: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 3$, $U = 4.00$, and $N^{max} = 1$.




Figure 249: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 3$, $U = 3.20$, and $N^{max} = 1$.




Figure 253: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 3$, $U = 4.00$, and $N^{max} = 2$.




Figure 254: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 3$, $U = 4.00$, and $N^{max} = 3$.




Figure 258: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 4$, $U = 3.20$, and $N^{max} = 1$.




Figure 255: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 3$, $U = 5.60$, and $N^{max} = 1$.




Figure 259: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 4$, $U = 3.20$, and $N^{max} = 2$.




Figure 256: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 3$, $U = 5.60$, and $N^{max} = 2$.




Figure 260: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 4$, $U = 3.20$, and $N^{max} = 3$.




Figure 257: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 3$, $U = 5.60$, and $N^{max} = 3$.




Figure 261: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 4$, $U = 4.00$, and $N^{max} = 1$.




Figure 262: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 4$, $U = 4.00$, and $N^{max} = 2$.




Figure 266: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 4$, $U = 5.60$, and $N^{max} = 3$.




Figure 263: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 4$, $U = 4.00$, and $N^{max} = 3$.




Figure 267: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 5$, $U = 3.20$, and $N^{max} = 1$.




Figure 264: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 4$, $U = 5.60$, and $N^{max} = 1$.




Figure 268: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 5$, $U = 3.20$, and $N^{max} = 2$.




Figure 265: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 4$, $U = 5.60$, and $N^{max} = 2$.




Figure 269: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 5$, $U = 3.20$, and $N^{max} = 3$.




Figure 270: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 5$, $U = 4.00$, and $N^{max} = 1$.




Figure 274: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 5$, $U = 5.60$, and $N^{max} = 2$.




Figure 271: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 5$, $U = 4.00$, and $N^{max} = 2$.




Figure 275: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 5$, $U = 5.60$, and $N^{max} = 3$.




Figure 272: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 5$, $U = 4.00$, and $N^{max} = 3$.




Figure 276: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 6$, $U = 3.20$, and $N^{max} = 1$.




Figure 273: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 5$, $U = 5.60$, and $N^{max} = 1$.




Figure 277: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 6$, $U = 3.20$, and $N^{max} = 2$.




Figure 278: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 6$, $U = 3.20$, and $N^{max} = 3$.




Figure 282: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 6$, $U = 5.60$, and $N^{max} = 1$.




Figure 279: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 6$, $U = 4.00$, and $N^{max} = 1$.




Figure 283: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 6$, $U = 5.60$, and $N^{max} = 2$.




Figure 280: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 6$, $U = 4.00$, and $N^{max} = 2$.




Figure 284: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 6$, $U = 5.60$, and $N^{max} = 3$.




Figure 281: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 6$, $U = 4.00$, and $N^{max} = 3$.




Figure 285: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 7$, $U = 3.20$, and $N^{max} = 1$.




Figure 286: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 7$, $U = 3.20$, and $N^{max} = 2$.




Figure 290: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 7$, $U = 4.00$, and $N^{max} = 3$.




Figure 287: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 7$, $U = 3.20$, and $N^{max} = 3$.




Figure 291: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 7$, $U = 5.60$, and $N^{max} = 1$.




Figure 288: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 7$, $U = 4.00$, and $N^{max} = 1$.




Figure 292: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 7$, $U = 5.60$, and $N^{max} = 2$.




Figure 289: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 7$, $U = 4.00$, and $N^{max} = 2$.




Figure 293: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 7$, $U = 5.60$, and $N^{max} = 3$.




Figure 294: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 8$, $U = 3.20$, and $N^{max} = 1$.




Figure 298: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 8$, $U = 4.00$, and $N^{max} = 2$.




Figure 295: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 8$, $U = 3.20$, and $N^{max} = 2$.




Figure 299: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 8$, $U = 4.00$, and $N^{max} = 3$.




Figure 296: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 8$, $U = 3.20$, and $N^{max} = 3$.




Figure 300: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 8$, $U = 5.60$, and $N^{max} = 1$.




Figure 297: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 8$, $U = 4.00$, and $N^{max} = 1$.




Figure 301: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 8$, $U = 5.60$, and $N^{max} = 2$.




Figure 302: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 20$, $n^{lat} = 8$, $U = 5.60$, and $N^{max} = 3$.




Figure 306: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 0$, $U = 4.00$, and $N^{max} = 1$.




Figure 303: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 0$, $U = 3.20$, and $N^{max} = 1$.




Figure 307: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 0$, $U = 4.00$, and $N^{max} = 2$.




Figure 304: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 0$, $U = 3.20$, and $N^{max} = 2$.




Figure 308: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 0$, $U = 4.00$, and $N^{max} = 3$.




Figure 305: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 0$, $U = 3.20$, and $N^{max} = 3$.




Figure 309: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 0$, $U = 5.60$, and $N^{max} = 1$.




Figure 310: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 0$, $U = 5.60$, and $N^{max} = 2$.




Figure 314: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 1$, $U = 3.20$, and $N^{max} = 3$.




Figure 311: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 0$, $U = 5.60$, and $N^{max} = 3$.




Figure 315: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 1$, $U = 4.00$, and $N^{max} = 1$.




Figure 312: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 1$, $U = 3.20$, and $N^{max} = 1$.




Figure 316: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 1$, $U = 4.00$, and $N^{max} = 2$.




Figure 313: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 1$, $U = 3.20$, and $N^{max} = 2$.




Figure 317: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 1$, $U = 4.00$, and $N^{max} = 3$.




Figure 318: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 1$, $U = 5.60$, and $N^{max} = 1$.




Figure 322: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 2$, $U = 3.20$, and $N^{max} = 2$.




Figure 319: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 1$, $U = 5.60$, and $N^{max} = 2$.




Figure 323: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 2$, $U = 3.20$, and $N^{max} = 3$.




Figure 320: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 1$, $U = 5.60$, and $N^{max} = 3$.




Figure 324: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 2$, $U = 4.00$, and $N^{max} = 1$.




Figure 321: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 2$, $U = 3.20$, and $N^{max} = 1$.




Figure 325: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 2$, $U = 4.00$, and $N^{max} = 2$.




Figure 326: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 2$, $U = 4.00$, and $N^{max} = 3$.




Figure 330: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 3$, $U = 3.20$, and $N^{max} = 1$.




Figure 327: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 2$, $U = 5.60$, and $N^{max} = 1$.




Figure 331: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 3$, $U = 3.20$, and $N^{max} = 2$.




Figure 328: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 2$, $U = 5.60$, and $N^{max} = 2$.




Figure 332: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 3$, $U = 3.20$, and $N^{max} = 3$.




Figure 329: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 2$, $U = 5.60$, and $N^{max} = 3$.




Figure 333: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 3$, $U = 4.00$, and $N^{max} = 1$.




Figure 334: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 3$, $U = 4.00$, and $N^{max} = 2$.




Figure 338: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 3$, $U = 5.60$, and $N^{max} = 3$.




Figure 335: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 3$, $U = 4.00$, and $N^{max} = 3$.




Figure 339: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 4$, $U = 3.20$, and $N^{max} = 1$.




Figure 336: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 3$, $U = 5.60$, and $N^{max} = 1$.




Figure 340: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 4$, $U = 3.20$, and $N^{max} = 2$.




Figure 337: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 3$, $U = 5.60$, and $N^{max} = 2$.




Figure 341: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 4$, $U = 3.20$, and $N^{max} = 3$.




Figure 342: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 4$, $U = 4.00$, and $N^{max} = 1$.




Figure 346: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 4$, $U = 5.60$, and $N^{max} = 2$.




Figure 343: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 4$, $U = 4.00$, and $N^{max} = 2$.




Figure 347: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 4$, $U = 5.60$, and $N^{max} = 3$.




Figure 344: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 4$, $U = 4.00$, and $N^{max} = 3$.




Figure 348: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 5$, $U = 3.20$, and $N^{max} = 1$.




Figure 345: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 4$, $U = 5.60$, and $N^{max} = 1$.




Figure 349: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 5$, $U = 3.20$, and $N^{max} = 2$.




Figure 350: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 5$, $U = 3.20$, and $N^{max} = 3$.




Figure 354: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 5$, $U = 5.60$, and $N^{max} = 1$.




Figure 351: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 5$, $U = 4.00$, and $N^{max} = 1$.




Figure 355: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 5$, $U = 5.60$, and $N^{max} = 2$.




Figure 352: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 5$, $U = 4.00$, and $N^{max} = 2$.




Figure 356: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 5$, $U = 5.60$, and $N^{max} = 3$.




Figure 353: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 5$, $U = 4.00$, and $N^{max} = 3$.




Figure 357: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 6$, $U = 3.20$, and $N^{max} = 1$.




Figure 358: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 6$, $U = 3.20$, and $N^{max} = 2$.




Figure 362: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 6$, $U = 4.00$, and $N^{max} = 3$.




Figure 359: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 6$, $U = 3.20$, and $N^{max} = 3$.




Figure 363: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 6$, $U = 5.60$, and $N^{max} = 1$.




Figure 360: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 6$, $U = 4.00$, and $N^{max} = 1$.




Figure 364: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 6$, $U = 5.60$, and $N^{max} = 2$.




Figure 361: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 6$, $U = 4.00$, and $N^{max} = 2$.




Figure 365: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 6$, $U = 5.60$, and $N^{max} = 3$.




Figure 366: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 7$, $U = 3.20$, and $N^{max} = 1$.




Figure 370: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 7$, $U = 4.00$, and $N^{max} = 2$.




Figure 367: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 7$, $U = 3.20$, and $N^{max} = 2$.




Figure 371: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 7$, $U = 4.00$, and $N^{max} = 3$.




Figure 368: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 7$, $U = 3.20$, and $N^{max} = 3$.




Figure 372: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 7$, $U = 5.60$, and $N^{max} = 1$.




Figure 369: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 7$, $U = 4.00$, and $N^{max} = 1$.




Figure 373: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 7$, $U = 5.60$, and $N^{max} = 2$.




Figure 374: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 7$, $U = 5.60$, and $N^{max} = 3$.




Figure 378: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 8$, $U = 4.00$, and $N^{max} = 1$.




Figure 375: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 8$, $U = 3.20$, and $N^{max} = 1$.




Figure 379: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 8$, $U = 4.00$, and $N^{max} = 2$.




Figure 376: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 8$, $U = 3.20$, and $N^{max} = 2$.




Figure 380: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 8$, $U = 4.00$, and $N^{max} = 3$.




Figure 377: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 8$, $U = 3.20$, and $N^{max} = 3$.




Figure 381: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 8$, $U = 5.60$, and $N^{max} = 1$.




Figure 382: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 8$, $U = 5.60$, and $N^{max} = 2$.




Figure 386: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 0$, $U = 3.20$, and $N^{max} = 3$.




Figure 383: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 30$, $n^{lat} = 8$, $U = 5.60$, and $N^{max} = 3$.




Figure 387: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 0$, $U = 4.00$, and $N^{max} = 1$.




Figure 384: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 0$, $U = 3.20$, and $N^{max} = 1$.




Figure 388: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 0$, $U = 4.00$, and $N^{max} = 2$.




Figure 385: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 0$, $U = 3.20$, and $N^{max} = 2$.




Figure 389: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 0$, $U = 4.00$, and $N^{max} = 3$.




Figure 390: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 0$, $U = 5.60$, and $N^{max} = 1$.




Figure 394: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 1$, $U = 3.20$, and $N^{max} = 2$.




Figure 391: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 0$, $U = 5.60$, and $N^{max} = 2$.




Figure 395: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 1$, $U = 3.20$, and $N^{max} = 3$.




Figure 392: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 0$, $U = 5.60$, and $N^{max} = 3$.




Figure 396: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 1$, $U = 4.00$, and $N^{max} = 1$.




Figure 393: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 1$, $U = 3.20$, and $N^{max} = 1$.




Figure 397: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 1$, $U = 4.00$, and $N^{max} = 2$.




Figure 398: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 1$, $U = 4.00$, and $N^{max} = 3$.




Figure 402: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 2$, $U = 3.20$, and $N^{max} = 1$.




Figure 399: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 1$, $U = 5.60$, and $N^{max} = 1$.




Figure 403: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 2$, $U = 3.20$, and $N^{max} = 2$.




Figure 400: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 1$, $U = 5.60$, and $N^{max} = 2$.




Figure 404: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 2$, $U = 3.20$, and $N^{max} = 3$.




Figure 401: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 1$, $U = 5.60$, and $N^{max} = 3$.




Figure 405: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 2$, $U = 4.00$, and $N^{max} = 1$.




Figure 406: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 2$, $U = 4.00$, and $N^{max} = 2$.




Figure 410: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 2$, $U = 5.60$, and $N^{max} = 3$.




Figure 407: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 2$, $U = 4.00$, and $N^{max} = 3$.




Figure 411: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 3$, $U = 3.20$, and $N^{max} = 1$.




Figure 408: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 2$, $U = 5.60$, and $N^{max} = 1$.




Figure 412: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 3$, $U = 3.20$, and $N^{max} = 2$.




Figure 409: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 2$, $U = 5.60$, and $N^{max} = 2$.




Figure 413: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 3$, $U = 3.20$, and $N^{max} = 3$.




Figure 414: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 3$, $U = 4.00$, and $N^{max} = 1$.




Figure 418: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 3$, $U = 5.60$, and $N^{max} = 2$.




Figure 415: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 3$, $U = 4.00$, and $N^{max} = 2$.




Figure 419: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 3$, $U = 5.60$, and $N^{max} = 3$.




Figure 416: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 3$, $U = 4.00$, and $N^{max} = 3$.




Figure 420: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 4$, $U = 3.20$, and $N^{max} = 1$.




Figure 417: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 3$, $U = 5.60$, and $N^{max} = 1$.




Figure 421: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 4$, $U = 3.20$, and $N^{max} = 2$.




Figure 422: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 4$, $U = 3.20$, and $N^{max} = 3$.




Figure 426: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 4$, $U = 5.60$, and $N^{max} = 1$.




Figure 423: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 4$, $U = 4.00$, and $N^{max} = 1$.




Figure 427: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 4$, $U = 5.60$, and $N^{max} = 2$.




Figure 424: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 4$, $U = 4.00$, and $N^{max} = 2$.




Figure 428: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 4$, $U = 5.60$, and $N^{max} = 3$.




Figure 425: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 4$, $U = 4.00$, and $N^{max} = 3$.




Figure 429: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 5$, $U = 3.20$, and $N^{max} = 1$.




Figure 430: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 5$, $U = 3.20$, and $N^{max} = 2$.




Figure 434: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 5$, $U = 4.00$, and $N^{max} = 3$.




Figure 431: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 5$, $U = 3.20$, and $N^{max} = 3$.




Figure 435: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 5$, $U = 5.60$, and $N^{max} = 1$.




Figure 432: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 5$, $U = 4.00$, and $N^{max} = 1$.




Figure 436: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 5$, $U = 5.60$, and $N^{max} = 2$.




Figure 433: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 5$, $U = 4.00$, and $N^{max} = 2$.




Figure 437: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 5$, $U = 5.60$, and $N^{max} = 3$.




Figure 438: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 6$, $U = 3.20$, and $N^{max} = 1$.




Figure 442: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 6$, $U = 4.00$, and $N^{max} = 2$.




Figure 439: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 6$, $U = 3.20$, and $N^{max} = 2$.




Figure 443: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 6$, $U = 4.00$, and $N^{max} = 3$.




Figure 440: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 6$, $U = 3.20$, and $N^{max} = 3$.




Figure 444: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 6$, $U = 5.60$, and $N^{max} = 1$.




Figure 441: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 6$, $U = 4.00$, and $N^{max} = 1$.




Figure 445: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 6$, $U = 5.60$, and $N^{max} = 2$.




Figure 446: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 6$, $U = 5.60$, and $N^{max} = 3$.




Figure 450: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 7$, $U = 4.00$, and $N^{max} = 1$.




Figure 447: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 7$, $U = 3.20$, and $N^{max} = 1$.




Figure 451: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 7$, $U = 4.00$, and $N^{max} = 2$.




Figure 448: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 7$, $U = 3.20$, and $N^{max} = 2$.




Figure 452: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 7$, $U = 4.00$, and $N^{max} = 3$.




Figure 449: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 7$, $U = 3.20$, and $N^{max} = 3$.




Figure 453: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 7$, $U = 5.60$, and $N^{max} = 1$.




Figure 454: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 7$, $U = 5.60$, and $N^{max} = 2$.




Figure 458: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 8$, $U = 3.20$, and $N^{max} = 3$.




Figure 455: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 7$, $U = 5.60$, and $N^{max} = 3$.




Figure 459: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 8$, $U = 4.00$, and $N^{max} = 1$.




Figure 456: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 8$, $U = 3.20$, and $N^{max} = 1$.




Figure 460: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 8$, $U = 4.00$, and $N^{max} = 2$.




Figure 457: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 8$, $U = 3.20$, and $N^{max} = 2$.




Figure 461: Schedulability under the OMLP and the OMIP for $m = 8$, $n = 40$, $n^{lat} = 8$, $U = 4.00$, and $N^{max} = 3$.




Figure 462: Schedulability under the OMLP and the OMIP for $m = 8, n = 40, n^{lat} = 8, U = 5.60,$ and $N^{max} = 1.$




Figure 466: Schedulability under the OMLP and the OMIP for $m = 16, n = 20, n^{lat} = 0, U = 6.40,$ and $N^{max} = 2.$




Figure 463: Schedulability under the OMLP and the OMIP for $m = 8, n = 40, n^{lat} = 8, U = 5.60,$ and $N^{max} = 2.$




Figure 467: Schedulability under the OMLP and the OMIP for $m = 16, n = 20, n^{lat} = 0, U = 6.40,$ and $N^{max} = 3.$




Figure 464: Schedulability under the OMLP and the OMIP for $m = 8, n = 40, n^{lat} = 8, U = 5.60,$ and $N^{max} = 3.$




Figure 468: Schedulability under the OMLP and the OMIP for $m = 16, n = 20, n^{lat} = 0, U = 8.00,$ and $N^{max} = 1.$




Figure 465: Schedulability under the OMLP and the OMIP for $m = 16, n = 20, n^{lat} = 0, U = 6.40,$ and $N^{max} = 1.$




Figure 469: Schedulability under the OMLP and the OMIP for $m = 16, n = 20, n^{lat} = 0, U = 8.00,$ and $N^{max} = 2.$




Figure 470: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 0$, $U = 8.00$, and $N^{max} = 3$.




Figure 474: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 1$, $U = 6.40$, and $N^{max} = 1$.




Figure 471: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 0$, $U = 11.20$, and $N^{max} = 1$.




Figure 475: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 1$, $U = 6.40$, and $N^{max} = 2$.




Figure 472: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 0$, $U = 11.20$, and $N^{max} = 2$.




Figure 476: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 1$, $U = 6.40$, and $N^{max} = 3$.




Figure 473: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 0$, $U = 11.20$, and $N^{max} = 3$.




Figure 477: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 1$, $U = 8.00$, and $N^{max} = 1$.




Figure 478: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 1$, $U = 8.00$, and $N^{max} = 2$.




Figure 482: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 1$, $U = 11.20$, and $N^{max} = 3$.




Figure 479: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 1$, $U = 8.00$, and $N^{max} = 3$.




Figure 483: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 2$, $U = 6.40$, and $N^{max} = 1$.




Figure 480: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 1$, $U = 11.20$, and $N^{max} = 1$.




Figure 484: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 2$, $U = 6.40$, and $N^{max} = 2$.




Figure 481: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 1$, $U = 11.20$, and $N^{max} = 2$.




Figure 485: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 2$, $U = 6.40$, and $N^{max} = 3$.




Figure 486: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 2$, $U = 8.00$, and $N^{max} = 1$.




Figure 490: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 2$, $U = 11.20$, and $N^{max} = 2$.




Figure 487: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 2$, $U = 8.00$, and $N^{max} = 2$.




Figure 491: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 2$, $U = 11.20$, and $N^{max} = 3$.




Figure 488: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 2$, $U = 8.00$, and $N^{max} = 3$.




Figure 492: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 3$, $U = 6.40$, and $N^{max} = 1$.




Figure 489: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 2$, $U = 11.20$, and $N^{max} = 1$.




Figure 493: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 3$, $U = 6.40$, and $N^{max} = 2$.




Figure 494: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 3$, $U = 6.40$, and $N^{max} = 3$.




Figure 498: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 3$, $U = 11.20$, and $N^{max} = 1$.




Figure 495: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 3$, $U = 8.00$, and $N^{max} = 1$.




Figure 499: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 3$, $U = 11.20$, and $N^{max} = 2$.




Figure 496: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 3$, $U = 8.00$, and $N^{max} = 2$.




Figure 500: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 3$, $U = 11.20$, and $N^{max} = 3$.




Figure 497: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 3$, $U = 8.00$, and $N^{max} = 3$.




Figure 501: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 4$, $U = 6.40$, and $N^{max} = 1$.




Figure 502: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 4$, $U = 6.40$, and $N^{max} = 2$.




Figure 506: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 4$, $U = 8.00$, and $N^{max} = 3$.




Figure 503: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 4$, $U = 6.40$, and $N^{max} = 3$.




Figure 507: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 4$, $U = 11.20$, and $N^{max} = 1$.




Figure 504: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 4$, $U = 8.00$, and $N^{max} = 1$.




Figure 508: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 4$, $U = 11.20$, and $N^{max} = 2$.




Figure 505: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 4$, $U = 8.00$, and $N^{max} = 2$.




Figure 509: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 4$, $U = 11.20$, and $N^{max} = 3$.




Figure 510: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 5$, $U = 6.40$, and $N^{max} = 1$.




Figure 514: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 5$, $U = 8.00$, and $N^{max} = 2$.




Figure 511: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 5$, $U = 6.40$, and $N^{max} = 2$.




Figure 515: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 5$, $U = 8.00$, and $N^{max} = 3$.




Figure 512: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 5$, $U = 6.40$, and $N^{max} = 3$.




Figure 516: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 5$, $U = 11.20$, and $N^{max} = 1$.




Figure 513: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 5$, $U = 8.00$, and $N^{max} = 1$.




Figure 517: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 5$, $U = 11.20$, and $N^{max} = 2$.




Figure 518: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 5$, $U = 11.2$, and $N^{max} = 3$.




Figure 522: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 6$, $U = 8.0$, and $N^{max} = 1$.




Figure 519: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 6$, $U = 6.4$, and $N^{max} = 1$.




Figure 523: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 6$, $U = 8.0$, and $N^{max} = 2$.




Figure 520: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 6$, $U = 6.4$, and $N^{max} = 2$.




Figure 524: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 6$, $U = 8.0$, and $N^{max} = 3$.




Figure 521: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 6$, $U = 6.4$, and $N^{max} = 3$.




Figure 525: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 7$, $U = 6.4$, and $N^{max} = 1$.




Figure 526: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 7$, $U = 6.40$, and $N^{max} = 2$.




Figure 530: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 7$, $U = 8.00$, and $N^{max} = 3$.




Figure 527: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 7$, $U = 6.40$, and $N^{max} = 3$.




Figure 531: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 0$, $U = 6.40$, and $N^{max} = 1$.




Figure 528: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 7$, $U = 8.00$, and $N^{max} = 1$.




Figure 532: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 0$, $U = 6.40$, and $N^{max} = 2$.




Figure 529: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 20$, $n^{lat} = 7$, $U = 8.00$, and $N^{max} = 2$.




Figure 533: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 0$, $U = 6.40$, and $N^{max} = 3$.




Figure 534: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 0$, $U = 8.00$, and $N^{max} = 1$.




Figure 538: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 0$, $U = 11.20$, and $N^{max} = 2$.




Figure 535: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 0$, $U = 8.00$, and $N^{max} = 2$.




Figure 539: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 0$, $U = 11.20$, and $N^{max} = 3$.




Figure 536: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 0$, $U = 8.00$, and $N^{max} = 3$.




Figure 540: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 1$, $U = 6.40$, and $N^{max} = 1$.




Figure 537: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 0$, $U = 11.20$, and $N^{max} = 1$.




Figure 541: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 1$, $U = 6.40$, and $N^{max} = 2$.




Figure 542: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 1$, $U = 6.40$, and $N^{max} = 3$.




Figure 546: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 1$, $U = 11.20$, and $N^{max} = 1$.




Figure 543: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 1$, $U = 8.00$, and $N^{max} = 1$.




Figure 547: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 1$, $U = 11.20$, and $N^{max} = 2$.




Figure 544: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 1$, $U = 8.00$, and $N^{max} = 2$.




Figure 548: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 1$, $U = 11.20$, and $N^{max} = 3$.




Figure 545: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 1$, $U = 8.00$, and $N^{max} = 3$.




Figure 549: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 2$, $U = 6.40$, and $N^{max} = 1$.




Figure 550: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 2$, $U = 6.40$, and $N^{max} = 2$.




Figure 554: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 2$, $U = 8.00$, and $N^{max} = 3$.




Figure 551: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 2$, $U = 6.40$, and $N^{max} = 3$.




Figure 555: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 2$, $U = 11.20$, and $N^{max} = 1$.




Figure 552: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 2$, $U = 8.00$, and $N^{max} = 1$.




Figure 556: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 2$, $U = 11.20$, and $N^{max} = 2$.




Figure 553: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 2$, $U = 8.00$, and $N^{max} = 2$.




Figure 557: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 2$, $U = 11.20$, and $N^{max} = 3$.




Figure 558: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 3$, $U = 6.40$, and $N^{max} = 1$.




Figure 562: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 3$, $U = 8.00$, and $N^{max} = 2$.




Figure 559: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 3$, $U = 6.40$, and $N^{max} = 2$.




Figure 563: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 3$, $U = 8.00$, and $N^{max} = 3$.




Figure 560: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 3$, $U = 6.40$, and $N^{max} = 3$.




Figure 564: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 3$, $U = 11.20$, and $N^{max} = 1$.




Figure 561: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 3$, $U = 8.00$, and $N^{max} = 1$.




Figure 565: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 3$, $U = 11.20$, and $N^{max} = 2$.




Figure 566: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 3$, $U = 11.20$, and $N^{max} = 3$.




Figure 570: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 4$, $U = 8.00$, and $N^{max} = 1$.




Figure 567: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 4$, $U = 6.40$, and $N^{max} = 1$.




Figure 571: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 4$, $U = 8.00$, and $N^{max} = 2$.




Figure 568: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 4$, $U = 6.40$, and $N^{max} = 2$.




Figure 572: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 4$, $U = 8.00$, and $N^{max} = 3$.




Figure 569: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 4$, $U = 6.40$, and $N^{max} = 3$.




Figure 573: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 4$, $U = 11.20$, and $N^{max} = 1$.




Figure 574: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 4$, $U = 11.2$, and $N^{max} = 2$.




Figure 578: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 5$, $U = 6.4$, and $N^{max} = 3$.




Figure 575: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 4$, $U = 11.2$, and $N^{max} = 3$.




Figure 579: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 5$, $U = 8.0$, and $N^{max} = 1$.




Figure 576: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 5$, $U = 6.4$, and $N^{max} = 1$.




Figure 580: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 5$, $U = 8.0$, and $N^{max} = 2$.




Figure 577: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 5$, $U = 6.4$, and $N^{max} = 2$.




Figure 581: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 5$, $U = 8.0$, and $N^{max} = 3$.




Figure 582: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 5$, $U = 11.20$, and $N^{max} = 1$.




Figure 586: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 6$, $U = 6.40$, and $N^{max} = 2$.




Figure 583: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 5$, $U = 11.20$, and $N^{max} = 2$.




Figure 587: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 6$, $U = 6.40$, and $N^{max} = 3$.




Figure 584: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 5$, $U = 11.20$, and $N^{max} = 3$.




Figure 588: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 6$, $U = 8.00$, and $N^{max} = 1$.




Figure 585: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 6$, $U = 6.40$, and $N^{max} = 1$.




Figure 589: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 6$, $U = 8.00$, and $N^{max} = 2$.




Figure 590: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 6$, $U = 8.00$, and $N^{max} = 3$.




Figure 594: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 7$, $U = 6.40$, and $N^{max} = 1$.




Figure 591: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 6$, $U = 11.20$, and $N^{max} = 1$.




Figure 595: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 7$, $U = 6.40$, and $N^{max} = 2$.




Figure 592: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 6$, $U = 11.20$, and $N^{max} = 2$.




Figure 596: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 7$, $U = 6.40$, and $N^{max} = 3$.




Figure 593: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 6$, $U = 11.20$, and $N^{max} = 3$.




Figure 597: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 7$, $U = 8.00$, and $N^{max} = 1$.




Figure 598: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 7$, $U = 8.00$, and $N^{max} = 2$.




Figure 602: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 7$, $U = 11.20$, and $N^{max} = 3$.




Figure 599: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 7$, $U = 8.00$, and $N^{max} = 3$.




Figure 603: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 0$, $U = 6.40$, and $N^{max} = 1$.




Figure 600: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 7$, $U = 11.20$, and $N^{max} = 1$.




Figure 604: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 0$, $U = 6.40$, and $N^{max} = 2$.




Figure 601: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 30$, $n^{lat} = 7$, $U = 11.20$, and $N^{max} = 2$.




Figure 605: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 0$, $U = 6.40$, and $N^{max} = 3$.




Figure 606: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 0$, $U = 8.00$, and $N^{max} = 1$.




Figure 610: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 0$, $U = 11.20$, and $N^{max} = 2$.




Figure 607: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 0$, $U = 8.00$, and $N^{max} = 2$.




Figure 611: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 0$, $U = 11.20$, and $N^{max} = 3$.




Figure 608: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 0$, $U = 8.00$, and $N^{max} = 3$.




Figure 612: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 1$, $U = 6.40$, and $N^{max} = 1$.




Figure 609: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 0$, $U = 11.20$, and $N^{max} = 1$.




Figure 613: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 1$, $U = 6.40$, and $N^{max} = 2$.




Figure 614: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 1$, $U = 6.40$, and $N^{max} = 3$.




Figure 618: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 1$, $U = 11.20$, and $N^{max} = 1$.




Figure 615: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 1$, $U = 8.00$, and $N^{max} = 1$.




Figure 619: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 1$, $U = 11.20$, and $N^{max} = 2$.




Figure 616: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 1$, $U = 8.00$, and $N^{max} = 2$.




Figure 620: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 1$, $U = 11.20$, and $N^{max} = 3$.




Figure 617: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 1$, $U = 8.00$, and $N^{max} = 3$.




Figure 621: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 2$, $U = 6.40$, and $N^{max} = 1$.




Figure 622: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 2$, $U = 6.40$, and $N^{max} = 2$.




Figure 626: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 2$, $U = 8.00$, and $N^{max} = 3$.




Figure 623: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 2$, $U = 6.40$, and $N^{max} = 3$.




Figure 627: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 2$, $U = 11.20$, and $N^{max} = 1$.




Figure 624: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 2$, $U = 8.00$, and $N^{max} = 1$.




Figure 628: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 2$, $U = 11.20$, and $N^{max} = 2$.




Figure 625: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 2$, $U = 8.00$, and $N^{max} = 2$.




Figure 629: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 2$, $U = 11.20$, and $N^{max} = 3$.




Figure 630: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 3$, $U = 6.40$, and $N^{max} = 1$.




Figure 634: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 3$, $U = 8.00$, and $N^{max} = 2$.




Figure 631: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 3$, $U = 6.40$, and $N^{max} = 2$.




Figure 635: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 3$, $U = 8.00$, and $N^{max} = 3$.




Figure 632: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 3$, $U = 6.40$, and $N^{max} = 3$.




Figure 636: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 3$, $U = 11.20$, and $N^{max} = 1$.




Figure 633: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 3$, $U = 8.00$, and $N^{max} = 1$.




Figure 637: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 3$, $U = 11.20$, and $N^{max} = 2$.




Figure 638: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 3$, $U = 11.20$, and $N^{max} = 3$.




Figure 642: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 4$, $U = 8.00$, and $N^{max} = 1$.




Figure 639: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 4$, $U = 6.40$, and $N^{max} = 1$.




Figure 643: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 4$, $U = 8.00$, and $N^{max} = 2$.




Figure 640: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 4$, $U = 6.40$, and $N^{max} = 2$.




Figure 644: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 4$, $U = 8.00$, and $N^{max} = 3$.




Figure 641: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 4$, $U = 6.40$, and $N^{max} = 3$.




Figure 645: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 4$, $U = 11.20$, and $N^{max} = 1$.




Figure 646: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 4$, $U = 11.20$, and $N^{max} = 2$.




Figure 650: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 5$, $U = 6.40$, and $N^{max} = 3$.




Figure 647: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 4$, $U = 11.20$, and $N^{max} = 3$.




Figure 651: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 5$, $U = 8.00$, and $N^{max} = 1$.




Figure 648: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 5$, $U = 6.40$, and $N^{max} = 1$.




Figure 652: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 5$, $U = 8.00$, and $N^{max} = 2$.




Figure 649: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 5$, $U = 6.40$, and $N^{max} = 2$.




Figure 653: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 5$, $U = 8.00$, and $N^{max} = 3$.




Figure 654: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 5$, $U = 11.20$, and $N^{max} = 1$.




Figure 658: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 6$, $U = 6.40$, and $N^{max} = 2$.




Figure 655: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 5$, $U = 11.20$, and $N^{max} = 2$.




Figure 659: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 6$, $U = 6.40$, and $N^{max} = 3$.




Figure 656: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 5$, $U = 11.20$, and $N^{max} = 3$.




Figure 660: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 6$, $U = 8.00$, and $N^{max} = 1$.




Figure 657: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 6$, $U = 6.40$, and $N^{max} = 1$.




Figure 661: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 6$, $U = 8.00$, and $N^{max} = 2$.




Figure 662: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 6$, $U = 8.00$, and $N^{max} = 3$.




Figure 666: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 7$, $U = 6.40$, and $N^{max} = 1$.




Figure 663: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 6$, $U = 11.20$, and $N^{max} = 1$.




Figure 667: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 7$, $U = 6.40$, and $N^{max} = 2$.




Figure 664: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 6$, $U = 11.20$, and $N^{max} = 2$.




Figure 668: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 7$, $U = 6.40$, and $N^{max} = 3$.




Figure 665: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 6$, $U = 11.20$, and $N^{max} = 3$.




Figure 669: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 7$, $U = 8.00$, and $N^{max} = 1$.




Figure 670: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 7$, $U = 8.00$, and $N^{max} = 2$.




Figure 674: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 7$, $U = 11.20$, and $N^{max} = 3$.




Figure 671: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 7$, $U = 8.00$, and $N^{max} = 3$.




Figure 675: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 16$, $U = 6.40$, and $N^{max} = 1$.




Figure 672: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 7$, $U = 11.20$, and $N^{max} = 1$.




Figure 676: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 16$, $U = 6.40$, and $N^{max} = 2$.




Figure 673: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 7$, $U = 11.20$, and $N^{max} = 2$.




Figure 677: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 16$, $U = 6.40$, and $N^{max} = 3$.




Figure 678: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 16$, $U = 8.00$, and $N^{max} = 1$.




Figure 682: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 16$, $U = 11.20$, and $N^{max} = 2$.




Figure 679: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 16$, $U = 8.00$, and $N^{max} = 2$.




Figure 680: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 16$, $U = 8.00$, and $N^{max} = 3$.




Figure 683: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 16$, $U = 11.20$, and $N^{max} = 3$.




Figure 681: Schedulability under the OMLP and the OMIP for $m = 16$, $n = 40$, $n^{lat} = 16$, $U = 11.20$, and $N^{max} = 1$.