Scaling Global Scheduling with Message Passing

Felipe Cerqueira Manohar Vanga Björn B. Brandenburg Max Planck Institute for Software Systems (MPI-SWS)

Abstract—Global real-time schedulers have earned the reputation of scaling poorly due to the high runtime overheads involved in global state management. In this paper, two mature implementations, one using fine-grained locking (SCHED_DEADLINE) and one using coarse-grained locking (LITMUS^{RT}'s G-EDF plugin), are evaluated and it is shown that, regardless of locking granularity, indeed neither scales well w.r.t. worst-case overheads due to excessive lock contention. To demonstrate that this is not an inherent limitation of global scheduling, the design of G-EDF-MP is presented, a global scheduler that uses message passing to avoid lock contention and cache-line sharing. It is shown to offer up to a 23- to 36-fold reduction in worst-case scheduling overhead on a 64core platform, which translates into much improved schedulability (in some cases, more than 120 additional tasks can be supported).

I. INTRODUCTION

Global real-time schedulers dispatch tasks to processors dynamically at runtime based on the state of all processors and all tasks. Such a global view of the system enables a wealth of attractive algorithmic properties—most famously optimality (*e.g.*, [6, 23, 36])—that have spurred considerable interest in global scheduling in recent years (*e.g.*, [2, 5, 10, 25, 27, 30], see [22] for a recent overview). However, in practice, the Achilles' Heel of global scheduling is runtime overheads [7, 13, 16, 18]: maintaining the required consistent global state can be costly, and becomes increasingly problematic with rising core counts. One might thus reasonably wonder: is it at all possible to *scale* global scheduling beyond a handful of cores?

In this paper, we answer this question in the affirmative. We first study two mature, radically different implementations of global real-time scheduling—LITMUS^{RT}'s *global earliest-deadline first* (G-EDF) scheduler [13, 17] and the SCHED_DEADLINE scheduler for stock Linux [26, 31, 32] (SD hereafter)—and show that they indeed do *not* scale. LITMUS^{RT}'s implementation uses a single *coarse-grained* lock to protect all scheduler state; its lack of scalability is thus expected. However, to our surprise, we found that SD, which uses *fine-grained* per-processor locks, does not scale either: it achieves much lower average-case overheads than LITMUS^{RT}, but only at the expense of crippling worst-case overheads.

Based on the observation that the main scalability bottleneck is not locking granularity *per se*, but rather the number of processors that may access the same lock(s), we then sidestep the global state management problem with a novel design based on *message passing*. In our new design, processors report state changes to a dedicated scheduling processor that is responsible for *all* scheduling decisions. As a result, the global state is not shared, scheduling decisions are implicitly serialized, and no scheduler lock is ever accessed by more than two processors, a potent combination that scales well: on a 64-core Intel Xeon platform, we observed a 23x (resp., 36x) reduction in worst-case scheduling overhead w.r.t. LITMUS^{RT} (resp., SD).

A. Motivation: The Case for Global Scheduling

The defining property of a global scheduler is that each task may execute on potentially any processor, depending on availability and each task's current priority. Implicit in this definition is the reason why it is difficult to implement global scheduling: to dynamically dispatch tasks according to processor availability, and in accordance with the scheduling policy (*e.g.*, in order of increasing deadlines under G-EDF), *a globally consistent snapshot* of the entire system state is needed. As globally shared (and frequently changing) state is well known to be a scalability bottleneck, it is not surprising that global schedulers are inherently more afflicted by runtime overheads than partitioned schedulers (in which tasks are statically assigned to processors and all dispatching decisions can be made locally).

Nonetheless, there are good reasons to not give up on global scheduling. As already mentioned, (practical) optimal multiprocessor real-time schedulers are necessarily global. Further, even if optimality of the scheduler is a secondary concern (or none at all, as it is admittedly often the case in practice), global scheduling offers advantages that make it preferable for certain applications. For example, in *adaptive systems* (*i.e.*, if task requirements change at runtime in reaction to environmental changes) and in *open systems* (*i.e.*, if tasks may exit or join the system dynamically), global scheduling is much easier to deal with precisely because it dispatches tasks dynamically and hence avoids the need for complex task mapping adjustments [12].

Other reasons speaking for global scheduling are that it is less susceptible to deadline misses due to intermittent overloads (as the load is spread across all processors) [22], that it offers lower average-case scheduling latencies (because no processor idles while any task is waiting) [22], that it is best suited for race-toidle energy conservation strategies [30], and that it permits much simpler locking mechanisms (*e.g.*, classic priority inheritance works as expected under global scheduling, but is ineffective under partitioned scheduling, see *e.g.* [13]).

To summarize, while global scheduling is not necessarily the best approach for *all* real-time workloads, there exist many applications that can benefit significantly from *efficient* OS support for global real-time scheduling for reasons of flexibility, optimality, and simplicity. Case in point: while Linux, QNX, and many other commercial real-time OSs offer flexible scheduling APIs that allow specifying arbitrary processor affinities [28], they *default* to global scheduling.

B. Focus of This Paper

Motivated by the overhead challenges surrounding global scheduling, our objective is to understand how global scheduling policies may be efficiently supported on multicore platforms that are large by today's standards.

The conference version of this extended tech report appears in the *Proceedings of the 20th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS 2014)*, April 2014, IEEE.

To this end, we study two mature and actively maintained open-source versions of G-EDF in Linux as reference implementations reflecting the current state of the art. We focus on the G-EDF policy in this paper for a number of reasons. For one, it has received considerable attention in prior work [2, 5, 10, 22, 27] and multiple independent implementations are readily available. Further, it is well-suited for the (predominantly) soft real-time workloads likely to be encountered on large multicore platforms because it is optimal (in the sense that it is not subject to utilization loss) if bounded deadline tardiness is permitted [23]. And finally, it is representative of the larger class of job-level fixed-priority schedulers, which includes both fixed-priority scheduling (the default in most commercial real-time OSs such as QNX and VxWorks) and other interesting (soft) real-time schedulers [25]. In short, if it is possible to scale G-EDF, then many other interesting policies can be implemented similarly.

Contributions. The primary contribution of this paper is to demonstrate that it is indeed possible to construct scalable global schedulers: global scheduling is not *inherently* limited by kernellevel lock and cache contention, though it is the case with many existing implementations. To support this claim:

- We first review the challenges inherent in global scheduling and the designs of SD and LITMUS^{RT}'s G-EDF plugin, and identify key performance and correctness limitations in Sec. II.
- 2) To overcome the scalability limitations of the current lockbased solutions, we introduce a novel scheduler design in Sec. III that is primarily based on message passing, and in which the remaining locks are accessed in pairwise fashion only.
- 3) We demonstrate that our design performs significantly better at scale than either of the two major current G-EDF implementations by means of average- and worst-case overhead measurements across six multicore platforms with 8, 16, 24, 32, and 64 cores in Sec. IV. For instance, in the 64-core setup, the worst-case scheduling overhead is reduced by a factor of 23 compared to LITMUS^{RT}'s G-EDF plugin, and by a factor of 36 compared to SD.
- 4) Finally, in Sec. V, we report on large-scale schedulability experiments that show that the reduction in overheads translates into significantly improved schedulability: when considering soft real-time schedulability (*i.e.*, bounded tardiness) in the 64-core setup, our message-passing-based implementation can support more than 120 additional realtime tasks than the best prior implementation.

A secondary contribution of this paper is the first direct comparison of LITMUS^{RT} and SD on the same hardware platform and with identical workloads, which is long overdue given their similarity in aim and differences in design.

We review related work in Sec. VI, and begin with a discussion of the scalability challenges posed by global scheduling.

II. SCALABILITY CHALLENGES

Before discussing the scalability limitations of existing implementations, we need to clarify the relevant metrics. In this paper, we focus on the scalability of *scheduling overheads*: the costs of making a scheduling decision and dispatching a task (discussed in detail in Sec. IV), which should only grow slowly (or ideally not at all) with increasing core counts.

Another source of overheads that affects tasks under global scheduling are *cache-related preemption and migration delays* (CPMD), which tasks incur due to a loss of cache affinity when they resume execution after a preemption or migration. While CPMD is a non-negligible source of overhead in practice (and accounted for in the experiments in Sec. V), it is less relevant in the context of this paper because CPMD is primarily hardware-dependent (*i.e.*, not affected by the kernel's scheduler implementation). Further, CPMD is not specific to global scheduling, but also affects partitioned (and semi-partitioned) scheduling, whereas overhead scalability is a challenge specific to global scheduling, scheduling (partitioning is trivially scalable w.r.t. overheads).

In the following, we let *m* denote the number of processors, and assume familiarity with the sporadic task model on behalf of the reader. In short, a *sporadic task* releases a sequence of *jobs*, where each job corresponds to an invocation of the task and has an associated release time and deadline. From the point of view of the Linux kernel, a sporadic task is simply a process that more or less regularly becomes ready for scheduling. We therefore use the terms 'task' and 'job' interchangeably when discussing implementation matters. With these clarifications in place, we next discuss why it is difficult for global schedulers to maintain low overheads with increasing processor counts.

A. Requirements and Challenges

Implementing a correct global scheduler requires the maintenance or discovery of a consistent global snapshot of the system that includes (i) the current task-to-processor mapping (*i.e.*, which tasks are scheduled) and (ii) the set of jobs that are eligible to run but not scheduled (to determine whether a preemption is required). Such a consistent view is necessary to avoid *priority inversions*, that is, to ensure that the *global scheduling invariant* is satisfied at all times: in a system with m processors and r ready tasks, the min(m, r) highest-priority tasks must be scheduled, where each task's current priority is determined by a policy such as EDF.

Given a consistent snapshot, making a correct scheduling decision (*i.e.*, one that maintains the global scheduling invariant) is trivial. However, since scheduling events-job arrivals, completions, suspensions, and resumptions-may occur concurrently on any processor, the state on which a scheduling decision is based may be invalidated before the scheduler has finished making its decision. There are two possible ways to deal with this issue: serialize all scheduling events (as done in LITMUS^{RT}), thereby avoiding concurrent state changes, or discard the invalidated decision and retry immediately (Linux and SD use this technique). Unfortunately, both approaches are problematic from a scalability perspective. Whereas the former approach, if implemented with a global lock, introduces an obvious contention hotspot, the latter can be problematic because concurrent state changes may race repeatedly (i.e., a processor may have to retry many times before arriving at a consistent decision). Worse, both contention and the likelihood of concurrent state changes increase with the number of cores. In fact, this very issue—maintaining a task-to-processor assignment that respects the global scheduling invariant in the face of concurrent scheduling events—is the root of most complexity and scalability issues in global schedulers.

Another, more analytical problem is posed by job release interrupts. Jobs are typically released (i.e., in Linux, processes become ready) in response to timer or device interrupts. Such interrupts are thus an inherent aspect of global real-time scheduling; however, they are not controlled by the scheduler. Rather, interrupts are dispatched in hardware and always take precedence over the OS's scheduling policy. This deviation from the scheduling policy causes considerable analytical complications. Since most published schedulability tests do not account for interrupts (e.g., see [22]), their delaying effects must be accounted for separately prior to normal schedulability analysis [19]. However, because it is in general not possible to predict on which processor(s) a task will execute, it must be assumed that, in the worst case, a job is delayed by all interrupts that occur while it is pending [13, 19], which introduces considerable pessimism. Minimizing the impact of job releases on already scheduled jobs is thus an important goal in the implementation of global schedulers [16].

To motivate the scalability problem further, and to illustrate the impact of the just-described key issues, we next review the two considered reference implementations and demonstrate their lack of scalability on a 64-core Intel Xeon platform.

B. The Design and Scalability of G-EDF in LITMUS^{RT}

LITMUS^{RT} [1], a real-time extension of the Linux kernel originally developed at UNC Chapel Hill, has included a G-EDF plugin since its initial version [20]. The version considered herein is called GSN-EDF because it was designed to support suspension-based locking protocols and <u>n</u>on-preemptable sections with O(1) priority inversions [11, 13]. GSN-EDF has been the main G-EDF implementation of LITMUS^{RT} since 2007 and has remained structurally unchanged since then. It has also been subject of several prior studies [7, 13, 16, 18, 20], which we revisit in the context of this work in Sec. VI.

The design of GSN-EDF is straightforward and favors worstcase predictability and conceptual simplicity over average-case optimizations: a *single* global spin lock protects both the shared ready queue (a binomial heap) and the processor mapping that links tasks to processors, and also serializes all scheduling decisions. This simple, coarse-grained synchronization approach has the advantage that it is relatively easy to reason about correctness, but obviously comes at a scalability tradeoff.

GSN-EDF's lack of scalability is readily apparent in Fig. 1, which shows an excerpt of the evaluation discussed in detail in Sec. IV. In short, we measured the average and maximum runtime overheads of GSN-EDF when managing 8, 16, 24, 32, 48, and 64 cores, respectively, of a 2 GHz 64-core Intel Xeon machine. Fig. 1 depicts the maximum and average costs of two key operations, namely the scheduling overhead (*i.e.*, assigning jobs to processors) and the release overhead (*i.e.*, adding jobs to the ready queue and triggering preemptions).

Both the maximum scheduling and release overheads exhibit a strong dependence on m and quickly become excessive beyond

16-24 cores, exceeding 1,000 μs and 600 μs , respectively, at 48 cores. Worse, this trend also manifests in the average-case overheads, which both reach about 500 μs at m = 64. Given that tasks with deadlines in the range of a few milliseconds are not uncommon, scheduling overheads of more than *one millisecond* are not viable and clearly show GSN-EDF's lack of scalability.

Prior studies using GSN-EDF [2, 10, 22, 27] used platforms with 4, 24, and 32 processors, a range in which GSN-EDF still exhibits overheads that are moderate in the context of Fig. 1. Nonetheless, in conjunction with the analytical limitations of current interrupt accounting techniques, the high release overheads proved problematic [16]. For this reason, GSN-EDF supports *dedicated interrupt handling* [16, 37], where a dedicated processor handles all release interrupts, and realtime tasks are scheduled only on the other m - 1 processors to shield them from interrupts. Given release overheads in the range of hundreds of microseconds (as evident in Fig. 1), this is a worthwhile tradeoff since the capacity lost to accounting-related pessimism far exceeds the "loss" of one processor [16].

C. The SCHED_DEADLINE Patch for Linux

The SD patch [26, 31, 32] for Linux was developed at the Scuola Superiore Sant'Anna in Pisa and fundamentally differs from GSN-EDF in both its design and its goal. While LITMUS^{RT} is primarily a prototyping framework and abstraction layer for experimental real-time research, SD is a comparably mature implementation that directly integrates with the kernel and aims at eventual inclusion into mainline Linux. Its design hence closely follows Linux's *push/pull* architecture, which is based on distributed, per-processor ready queues, each guarded by a separate spin lock, which reduces both lock and cache contention by minimizing data sharing across processors. GSN-EDF and SD thus represent two extreme opposites: while in GSN-EDF global coordination is the default (and all scheduling decisions are serialized), in the push/pull design, most scheduling decisions are made locally (and concurrently).

In Linux, a global policy such as G-EDF must be implemented on top of the per-processor ready queues by explicitly loadbalancing the highest-priority tasks whenever local scheduling events require a global adjustment (e.g. when new jobs are released or when jobs complete). This is accomplished with the eponymous *push* and *pull* operations, which enact sourceand *target-initiated* migrations, respectively, to maintain the global scheduling invariant.¹ Conceptually, pull operations scan all other ready queues to "steal" backlogged higher-priority tasks before a local scheduling decision is made, and push operations move local backlogged tasks to other processors where they can be scheduled immediately after a local scheduling decision was made. In the actual implementation, these operations are not invoked as part of every scheduler invocation; rather, processors push whenever a new job is released or a local job is preempted, and pull when a local job completes or suspends.

¹While the Linux real-time scheduler defaults to global scheduling, it supports specifying *arbitrary processor affinities* to restrict a task's execution to a subset of processors. This allows emulating partitioned, global, and hybrid schedulers [28]. We focus on global affinities herein, as those are most difficult to support—adding affinity restrictions makes the (runtime) problem easier.



Fig. 1: Scheduling and release overheads under GSN-EDF. The average scheduling and release overheads coincide at this resolution.



Fig. 2: Scheduling and release overheads under SD. The average scheduling and release overheads coincide at this resolution.

Both operations require a pairwise acquisition of the relevant per-processor locks in order to maintain consistency. To prevent deadlocks, the locks are always acquired in order of increasing processor indices, which may require processors to drop and reacquire their local scheduler lock. To improve blocking time in the average case, the local lock is always dropped in Linux and both source and target locks are reacquired in order.

In summary, the push/pull design underlying the SD scheduler strives to avoid scalability bottlenecks by limiting cache contention through the use of per-processor run queues, and by using fine-grained, per-processor locks. Consequently, our starting assumption was that SD would easily and consistently outperform GSN-EDF. Surprisingly, we found that the push/pull design works both really well and not at all: as shown in Fig. 2, it yields very low average-case overheads, but it is also subject to excessive worst-case overheads.

On the positive side, our experiments show that the more sophisticated locking mechanism is very useful for coping with contention in the average case. We found the average-case overheads to remain low throughout the tested range, and well below 100 μs even for m = 64. From a throughput perspective, SD thus performs excellently, and much better than GSN-EDF.

However, from a (hard) real-time perspective, the results are troublesome, as the maximum observed overheads clearly do not scale well. In fact, in some configurations, SD exhibits maximum overheads more than five times worse than the much simpler GSN-EDF: for instance, for m = 24, the maximum scheduling overhead of GSN-EDF is roughly 180 μs , whereas SD already exceeds 1,000 μs . In the extreme of m = 64, we even observed scheduling overheads approaching 4.5 ms, which demonstrates a total lack of scalability w.r.t. worst-case overheads.

D. Root Causes: Transitive Blocking and Push Failures

As we set out to investigate the scalability of G-EDF implementations, our initial assumption was that per-processor state and locks as employed by SD would "naturally" be the most promising approach, a fallacy that we believe to be widespread. In fact, to the best of our knowledge, our experiments are the first to demonstrate that, w.r.t. maximum overheads, the push/pull approach scales significantly worse than a single, coarse-grained lock. In the end, it turns out that "simulating" a unified global state from multiple distributed states requires complex synchronization logic (*e.g.*, scanning operations, the release and reacquisition of locks, complex interaction among processors, *etc.*), which is detrimental to worst-case overheads.

A key reason for the observed spikes is that fine-grained locking gives rise to pathological, difficult to anticipate blocking scenarios when multiple processors push and pull simultaneously. While the per-processor locks are *on average* accessed by only a handful of processors, it is *possible* for any lock to be contested by all processors at the same time. Further, due to the scanning operations and the need for pairwise ordered lock acquisitions described above, this effect can escalate due to *transitive blocking* when high-contention scenarios are encountered on every outer and every nested lock acquisition.

The case for push/pull-based designs is further weakened by an "optimization" inherited from stock Linux's push/pull scheduler that prevents SD from maintaining the global invariant at all times (*i.e.*, priority inversions, though likely infrequent, are possible with the current implementation). The problem arises because a push operation may fail as follows. To carry out a push migration, the source processor first searches for the processor that is currently executing the lowest-priority task (*i.e.*, latest-deadline task in SD) by iteratively examining all processor states. After the target processor has been identified, its lock is reacquired to enact the migration. However, in the mean time while the source processor was not holding the target processor's lock, the target processor's state may have changed, thus invalidating the decision. Specifically, after deciding to migrate a task T_i to a processor P_k , a new, higher-priority task may arrive on P_k , either because it resumes or because another processor pushes first (as described in further detail in [13, Ch. 3] in the context of the stock Linux fixed-priority push/pull scheduler). This poses a fundamental question: in a push/pull design, how often must a processor retry a failed push operation in the worst case?

In the current SD implementation (version V8), if a push operation fails three times, then it is abandoned and the task is kept in the local queue even if that violates the global scheduling invariant. However, a retry limit of three is clearly insufficient, since with tens of cores and hundreds of tasks, it is certainly possible that more than three higher-priority tasks arrive concurrently with a push operation. Unfortunately, the hard-coded retry limit exists for a reason: correcting it would likely cause a further increase in (worst-case) overheads due to additional contention from repeatedly failing push operations.

Finally, it should be noted that overheads were measured while executing simple, periodic, CPU-bound, synchronously² released real-time tasks, as explained in Appendix A. The

²In a *synchronous* task set release, all tasks release their first job at the same time, which ensures that all tasks share a common time zero. In our case, the synchronous release serves to provoke worst-case contention and overheads.

synthetic benchmark tasks used to evaluate SD correspond to the standard workload used to benchmark LITMUS^{RT} (see Sec. IV) and mimic well-behaved recurrent real-time applications such as process control or signal processing tasks. We thus believe this to be a fair and relevant assessment of the two G-EDF implementations, but note that "antagonistic" tasks could likely provoke even higher overheads under either scheduler. Nonetheless, the results allow us to identify the central scalability bottleneck.

E. Scalability Implications

In the end, we found that neither fine-grained nor coarsegrained locking scales particularly well. This is because the true bottleneck is not *lock granularity*, which is more relevant in a throughput context, but rather *peak contention*, that is, the number of processors that potentially access the *same* lock(s). Peak contention, however, cannot be reduced by splitting the shared global state into ever finer pieces that still must be shared among all processors. To the contrary, due to transitive blocking, m processors sharing m locks in a nested fashion suffer larger worst-case delays than m processors sharing a single lock.

Ideally, a scheduler should have the simplicity of GSN-EDF, provide average-case overheads close to those of SD, and provide worst-case overheads as predictable and simple to reason about as those of GSN-EDF (but preferably lower). In the following, we present a novel G-EDF implementation, based on message passing instead of state sharing, that comes much closer to these goals than any prior implementation known to us.

III. GLOBAL SCHEDULING VIA MESSAGE PASSING

The use of message passing is a growing trend in recent work on multicore scalability [9, 21, 33]. Its main advantage is that it makes sharing explicit and forces algorithm designs that inherently provide data isolation and reduce sharing. In the following, we show that message passing is also a viable implementation technique at the timescale of a real-time scheduler. We begin with an overview of G-EDF-MP's design, our new messagepassing-based G-EDF implementation, and provide a detailed description of how it operates in Sec. III-B.

A. Design Overview

As discussed in Sec. II, our solution is motivated by the fact that large overheads from peak contention and cache-line sharing arise because each processor must coordinate with other processors to generate a globally consistent snapshot of the system prior to making a decision.

G-EDF-MP avoids these two issues through a simple tradeoff. Taking dedicated interrupt handling as already implemented in GSN-EDF a significant step further, it declares a single *dedicated scheduling processor* (henceforth the DSP) to be responsible for all scheduling decisions and all device and timer interrupts. The rest of the *client* processors in the system simply carry out the decision made by the DSP. To avoid analytical complications, the DSP itself does not execute any real-time tasks, which execute only on the shielded client processors.

Since only the DSP makes decisions, it is the only processor that needs to have a global view of the system. G-EDF-MP thus maintains this view locally in the DSP: ready queue updates do



not require any locking, to the effect that job release interrupts cannot be delayed by lock contention.

However, for reasons of efficiency, G-EDF-MP is not a "pure" message-passing-based implementation, as it still uses some locks to synchronize the assignment of tasks to clients. In particular, the DSP-to-client communication is carried out with two lock-protected task pointers per client, which the DSP updates in order to indicate task assignments. However, the peak contention of these locks is limited because (**i**) they are never acquired in a nested fashion, and (**ii**) they are subject to only pairwise sharing (by the DSP and the particular client).

The DSP uses *inter-processor interrupts* (IPIs) to notify clients when a particular decision requires an immediate preemption. Similarly, a client communicates local scheduling events (*e.g.* job completions, suspensions, *etc.*) to the DSP by writing a message to one of a small number of shared, wait-free FIFO mailboxes and sending an IPI as a notification. This allows the DSP to keep track of the global state of the system and to react as clients become available for other work.

Overall, the G-EDF-MP approach is conceptually simple, but requires careful implementation choices, which we discuss next.

B. The G-EDF-MP Data Structures

Fig. 3 provides an overview of the system. The DSP uses its local *DSP state* to make scheduling decisions. The clients have a small *per-client* state that is protected by per-client locks (not shown), which indicates the assigned and scheduled tasks. Preemptions are enacted by the DSP by modifying a client's state and sending an IPI. Clients communicate with the DSP by writing to a *shared mailbox* and sending an IPI to the DSP.

Per-client state. G-EDF-MP is conceptually based on the linkbased scheduling approach first introduced with GSN-EDF [17].

Link-based scheduling tracks two task pointers for each processor, the linked task that *should be* scheduled, and the scheduled task that currently *is* scheduled. The DSP updates the linked task as required to maintain the global scheduling invariant, and the client is responsible for tracking linked with scheduled as quickly as possible by context-switching to linked whenever the DSP updates the assignment. (This delay is accounted for analytically in Appendix B.)

The two main advantages of link-based scheduling are that it enables the implementation of non-preemptable critical section support with O(1) priority inversion [13, 15] (though this is not yet activated in G-EDF-MP), and that it cleanly decouples the logical scheduling aspects as dictated by the scheduling policy from the peculiarities of hardware and process management [13].

Each client's linked and scheduled fields are protected by a spin lock. Only clients may change their scheduled



Fig. 4: Overview of G-EDF-MP operations.

pointer, but both the DSP and the clients update the linked pointers: clients invalidate their linked pointer to indicate job completions or suspensions, and the DSP updates the linked pointer to assign new work. In contrast, under GSN-EDF, all processors may relink any other processor and all linked and scheduled pointers are protected by a single lock.

DSP state. The DSP state consists of (i) a ready queue of tasks that are ready to execute but are waiting to be linked, (ii) binomial heaps and timers tracking future job releases,³ and (iii) a copy of each client's state (*e.g.* which task is linked to, whether it is idle *etc.*) that is updated whenever the DSP changes linked or when a client reports idleness. Part (iii) is used to detect when preemptions become necessary. The DSP uses a copy of the client state (rather than accessing it directly) to improve cache locality and to avoid having to lock per-client states during preemption checks, which are frequently required. Notably, the DSP state is not protected by any locks since it is accessed only by the DSP with disabled interrupts.

Shared mailbox. The mailbox enables asynchronous client-to-DSP communication, where each message consists of a pointer to a callback function, the sender ID, and the task that triggered the scheduling event. It consists of a small number of Feather-Trace buffers [17], which support wait-free writes so clients never block when sending messages.⁴ To take advantage of the cache topology, there is one mailbox per socket that is shared only by cores on the same die (and the DSP). The DSP services message buffers in a round-robin fashion to avoid starvation, and briefly enables interrupts between messages (lest latency spikes arise). We next describe the main scheduling events.

C. The G-EDF-MP Operations

There are six major message types in G-EDF-MP, as illustrated in Fig. 4, which shows a sequence of events and messages.

1) When a job is released or resumed by an external event, the DSP receives an interrupt (since we use dedicated interrupt handling) and the associated handler is executed. Otherwise, if a task is woken up by another task on a client, the DSP is notified by the client via a message. In either case, the DSP immediately links the task to an idle core if one is available, and otherwise checks whether a preemption is in order. Since the DSP has a

³LITMUS^{RT} provides binomial heaps to enable the release of many jobs at once with efficient $O(\log n)$ merge operations [18]. G-EDF-MP simply reuses the existing infrastructure, with the difference that no locking is required.

⁴We chose Feather-Trace buffers because they are readily available in LITMUS^{RT} as part of the overhead tracing infrastructure. In future work, it may be interesting to evaluate whether specific-purpose, single-writer, cache-coherency-friendly message buffers [9, 21] achieve even lower overheads.

copy of all current task assignments, these decisions are made purely locally without any lock acquisitions.

In Fig. 4, the DSP decides that the newly arrived task T_2 should preempt task T_1 on the depicted client processor. To enact the preemption, the DSP acquires the client's lock, updates its linked pointer, and sends an IPI. After releasing the client's lock, the DSP adds the unlinked task T_1 back into the ready queue. Note that T_1 remains scheduled until the client receives the IPI, which illustrates the separation of logical and physical processor management in link-based scheduling.

2) When a client processor receives a rescheduling IPI, it locks its local state and simply checks whether linked differs from scheduled, and if so atomically updates scheduled to indicate that it has picked up the link change. It then unlocks its local state and carries out a context switch to the new task.

3) When a job suspends on a client processor, it immediately becomes unavailable for execution. Hence, the client invalidates its linked field to indicate that it is idle and sends a message to the DSP to communicate its availability. Since the client at this point does not have information about which job to schedule next (or even whether more jobs are pending), it simply idles (or services non-real-time background tasks).

Introducing idle time of course introduces additional delays for any pending real-time tasks, which must be reflected when accounting for overheads, as we discuss in Appendix B. Note that this idle time is not in conflict with the fact that G-EDF is a work-conserving policy. In any practical scheduler implementation, processors incur brief periods of unavailability while selecting the next task to dispatch, regardless of whether this is realized with locking or message passing mechanisms.

4) When the DSP is notified of a suspension, it discards the suspending task from its data structures. The suspending task will be "reintroduced" to the scheduler when it is later resumed by an interrupt or another task—see (1) above. Note that a suspending task might already be queued in the ready queue if a preemption IPI sent from the DSP to the client raced with a suspension IPI sent from the client to the DSP. In general, since G-EDF-MP intentionally aims to decouple clients from the DSP as much as possible for scalability reasons, many potential race conditions among in-flight events must be handled.

Since a suspension message implies that the sending client became idle, the DSP links the highest-priority job in the ready queue to the sender (if any). In the example in Fig. 4, the DSP links T_1 to the client and sends a preemption IPI, which causes T_1 to be dispatched as discussed in (2) above.

5) When a job completes on a client processor, the client reacts as in (3): it invalidates linked and notifies the DSP.

6) When the DSP is notified of a completion, if the job completed before its next release time, the DSP programs a local timer for the next job release; otherwise, if the job was tardy, it adds the task immediately back into the ready queue. Either way, it selects the highest-priority job in the ready queue (if any) and assigns it to the sender.

Concurrency challenges. The preceding list of events covers the major scheduling operations in G-EDF-MP. While the design

is intentionally simple at the conceptual level, several corner cases must be correctly dealt with in practice. For example, since an unlinked task T_i in the ready queue may remain briefly scheduled on some processor P_k (until the client's "is"-state catches up with the "should"-state), it may be picked by the DSP and linked to some other processor P_l that scheduled some other task T_i . In such a case, P_l must delay its context switch from T_i to T_i until P_k ceases to use T_i 's kernel stack.

This, however, becomes problematic if subsequently T_i is assigned to P_k —a deadlock results, as P_l uses T_j 's kernel stack while it waits for P_k to release T_i 's kernel stack, and vice versa. If left unchecked, this can happen as a result of short suspensions (e.g., due to I/O), and must be detected and resolved by swapping the links of T_i and T_j . To avoid such situations altogether, G-EDF-MP performs link swaps whenever it links a still-scheduled task. To allow the DSP to reliably detect still-scheduled tasks, clients must check linked and update scheduled atomically, as discussed in operation (2).

While such corner cases add some implementation complexity, they do not affect the overheads and scalability of G-EDF-MP negatively, as we show next.

IV. EVALUATION

To assess the scalability of G-EDF-MP, SD, and GSN-EDF in terms of average and maximum overheads, we conducted extensive experiments on a 64-core Intel Xeon X7550 2.0 GHz platform using LITMUS^{RT} version 2013.1 and version V8 of the SD patch, both based on Linux 3.10. Features that lead to unpredictability such as hardware multithreading, frequency scaling, and deep sleep states were disabled for all kernels, along with every kernel configuration option associated with debugging. All non-essential background services such as cron were disabled. For each $m \in \{8, 16, 24, 32, 48, 64\}$, each of the three kernel versions was compiled with support for at most mprocessors and booted on the test platform.

Overheads were measured with Feather-Trace [17], a lightweight tracing framework included in LITMUS^{RT}. As Feather-Trace was not built with large core counts in mind, the stock version uses a single, shared trace buffer, which itself creates a scalability bottleneck at higher core counts. To resolve this, we changed Feather-Trace to record most samples in processor-local buffers. We also added support for coping with non-synchronized cycle counters by determining each processor's cycle counter offset prior to each experiment.

For each m and for each $n \in \{m, 2m, \dots, 10m\}$, we generated five task sets with n synthetic tasks using Emberson et al.'s task set generator [24]. Task sets were generated with a target utilization chosen uniformly at random from [0.4m, 0.8m]and periods drawn from a log-uniform distribution in the range [10ms, 100ms]. Each task set was executed for 30 s under each of the following four schedulers: G-EDF-MP, SD, GSN-EDF without dedicated interrupt handling, and GSN-EDF with dedicated interrupt handling (GSN-EDF-DI hereafter). Note that real-time tasks execute on only m-1 of the m cores under G-EDF-MP and GSN-EDF-DI. We ran a cache-thrashing background workload to provoke maximum memory hierarchy



contention. In total, we traced 1,200 task sets during ten hours of real-time execution and collected over 7 billion raw samples.

Next, we discuss the different runtime overheads that delay real-time tasks, which are all traced individually in LITMUS^{RT}.

A. Measured Runtime Overheads

In the case of G-EDF-MP, nine major overhead sources affect a task's response time, which are illustrated in Fig. 5.

First, a task incurs *release event latency* (Δ^{ev}) if its triggering interrupt is delayed because the DSP disabled interrupts. Release overhead (Δ^{rel}) reflects the cost of processing the release interrupt, including all ready queue updates, preemption checks, and link changes. After the link change, the DSP sends a rescheduling IPI, which incurs *IPI latency* (Δ^{ipi}) due to interconnect contention and when the recipient disabled interrupts.

On the client, the task incurs scheduling overhead (Δ^{sch}) while the kernel selects the next task to be dispatched; this overhead includes any blocking delays incurred when accessing the scheduler state. Finally, *context-switch overhead* (Δ^{cxs}) is incurred when switching kernel stacks. While executing, a realtime task can be delayed by the periodic timer tick (which fires once every millisecond), incurring *tick overhead* (Δ^{tck}).

Two overheads specific to G-EDF-MP are the IPI sent to notify the DSP of pending messages and the time spent by the DSP reacting to messages. Regarding the former, we distinguish between Δ^{ipi} , the latency of normal reschedule IPIs sent to trigger the scheduler on a remote processor, and client request *latency* (Δ^{req}), which reflects the delay from the time that a client sends a message to the DSP until the DSP dequeues that specific message. Regarding the latter G-EDF-MP-specific overhead, each message has a corresponding callback function, the invocation of which causes DSP handler overhead (Δ^{dsp}).

CPMD (Δ^{cpd}) affects a task when it reloads cache contents after a preemption or migration. CPMD is not affected by the choice of G-EDF implementation and was measured separately using the method proposed by Bastoni et al. [8].

With the exception of the client request latency and DSP handler overhead, all of the just-discussed overhead sources also affect GSN-EDF, GSN-EDF-DI, and SD. To obtain comparable measures for SD, we ported Feather-Trace to the SD kernel (but did not introduce any other changes). Further, the scheduling overhead Δ^{sch} is actually measured in two parts (before and after the context switch); we report the aggregate measure here. In contrast to earlier LITMUS^{RT}-based studies [7, 16, 18], we did not apply any statistical outlier filters (which made possible

by recent LITMUS^{RT} improvements [14]). Finally, it is also worth mentioning that synchronously releasing all tasks (as discussed in Sec. II-D and Appendix A) creates a challenging benchmark, as it forces all real-time tasks to be activated together. The overhead measurements that were obtained thus reflect very unfavorable scenarios in terms of contention.

All overhead graphs are available in Appendix C, and the corresponding histograms in Appendix D. Due to space constraints, we discuss only the most pertinent trends in the following.

B. Experimental Results

As already shown in Figs. 1 and 2, GSN-EDF does not scale well w.r.t. both average- and worst-case overheads, and SD does not scale well w.r.t. worst-case overheads. The key questions thus are, (i) does G-EDF-MP scale well with the number of cores and the number of tasks, and (ii) does G-EDF-MP achieve low average-case overheads similar to SD?

We begin with question (i), and the answer is clearly yes. Fig. 6(a) shows the maximum observed scheduling overhead as a function of m. While the maximum scheduling overhead under GSN-EDF and SD increases from 40.19 μs and 52.11 μs , resp., at 8 cores to 2759.04 μs and 4337.55 μs , resp., at 64 cores, the scheduling overhead under G-EDF-MP increases from 12.75 μs to only 119.37 μs ; a relatively modest increase and a 23x and 36x improvement over GSN-EDF and SD, resp.

Similar trends are observed with increasing n, as shown in Fig. 6(c), which depicts Δ^{sch} as a function of n for m = 48. Whereas the maximum scheduling overhead under GSN-EDF, GSN-EDF-DI, and SD either already start high or reveal large upwards slopes, G-EDF-MP maintains near-constant overheads around 60 μs for the entire range of tested task set sizes. These trends are not surprising, since GSN-EDF, GSN-EDF-DI, and SD must acquire one or more globally shared locks (for which contention increases with m) and manipulate ready queues (which grow in size with n) as part of each scheduler invocation. Further, overheads under SD are strongly affected by increasing *n* since a larger number of tasks increases the frequency of scheduler invocations, which increases the likelihood of push failures and transitive blocking. In contrast, each client state lock in G-EDF-MP is accessed by only two processors, and clients never access the ready queue—the client-side scheduler is O(1).

GSN-EDF and GSN-EDF-DI exhibit mostly similar trends, with overheads slightly lower under GSN-EDF-DI simply because one fewer core participates in the scheduling, thereby reducing peak contention slightly.

The major difference in scalability between the primarily lockbased schedulers and G-EDF-MP also manifests in the maximum release overhead as shown in Fig. 6(b). The overhead under G-EDF-MP exhibits quicker growth than in Fig. 6(a) due to $O(\log n)$ updates to the ready queue (a binomial heap) and O(m) preemption checks. Still, G-EDF-MP exhibits much lower maximum release overhead than its counterparts afflicted by lock contention: for m = 64, the maximum release overhead exceeds 1700 μs and 1500 μs under GSN-EDF and SD, resp., whereas G-EDF-MP stays well below 275 μs .

To answer question (ii), consider the plots of average-case scheduling and release overheads shown in Figs. 6(d) and 6(e),

resp. As discussed in Sec. II, GSN-EDF does not scale well in the average case, which also extends to GSN-EDF-DI, whereas SD has no problems scaling in the average case (since pathological contention is rare). G-EDF-MP in turn exhibits still lower scheduling overheads, as seen in Fig. 6(d) (at m = 64, $\Delta^{sch} = 3.04 \,\mu s$ under G-EDF-MP vs. $\Delta^{sch} = 48.48 \,\mu s$ under SD).

In Fig. 6(e), G-EDF-MP exhibits somewhat higher average release overheads than SD (at m = 64, $\Delta^{rel} = 109.40 \,\mu s$ under G-EDF-MP vs. $\Delta^{rel} = 65.67 \,\mu s$ under SD), but more importantly, the overhead under G-EDF-MP grows only slowly from m = 8 to m = 64, in contrast to GSN-EDF (which is below $6 \,\mu s$ at m = 8, but exceeds $475 \,\mu s$ at m = 64).⁵ We conclude that G-EDF-MP already performs reasonably well in terms of average-case overheads, but also believe that there is room for further low-level code optimizations in the future.

Although the trends are promising so far, there is no silver bullet for multicore scalability—the design of G-EDF-MP introduces new overheads that also need to be considered. These are shown in Fig. 6(f), which depicts the average and maximum client request latency and the DSP handler cost.

The DSP handler cost is not particularly problematic since it is dominated by ready queue and preemption checks and thus grows only slowly with m, just like the release overhead.

The client request latency is determined by three factors: how long the DSP disables interrupts, how expensive it is for the DSP to dequeue a message, and how many messages are already queued (recall that each message buffer is FIFO-ordered). To keep the first factor low, the DSP enables interrupts between any two messages. Concerning the second factor, reading the mailbox becomes more expensive with growing core counts as more processors send messages concurrently. However, since we use per-socket mailboxes, this increase is relatively small.

The third factor, the length of the message queue, is the main bottleneck of G-EDF-MP—with an increasing number of clients, the peak contention in terms of simultaneous message arrivals naturally grows as well. This can be clearly seen in Fig. 6(f), where the maximum client request latency grows quickly with m. The average client request latency, however, is not strongly affected by this—again, peak contention is a rare event.

Finally, to offer further support for our claims and to illustrate how overheads vary with increasing task counts, Fig. 7 depicts six additional graphs showing scheduling overheads, release overheads, DSP handler overheads, and client request latency as a function of n on the 32-core platform. The general trend is that G-EDF-MP performs well w.r.t. the overheads encountered by all schedulers—in insets (a), (b), (d), and (e), G-EDF-MP exhibits the lowest overheads of the four schedulers for any n. In fact, in special cases, G-EDF-MP achieves even up to two orders of magnitude lower overheads than the reference implementations: the maximum observed scheduling overhead under G-EDF-MP is $\approx 17 \,\mu s$ vs. $\approx 1740 \,\mu s$ under SD, as apparent in Fig. 7(a). Finally, Fig. 7(c) confirms that client request latency is the main bottleneck w.r.t. worst-case overheads, but not a major problem

⁵Release operations are costly in both GSN-EDF and GSN-EDF-DI because, in the worst case, all tasks are released simultaneously and every processor must be preempted, which requires many updates of the global state and results in significant lock and cache contention.



Fig. 6: Runtime overheads under GSN-EDF, GSN-EDF-DI, SD, and G-EDF-MP. Note that each graph uses a different scale due to the wide range of observed overhead magnitudes.

w.r.t. average-case overheads, and Fig. 7(f) shows that the cost per DSP operation is virtually constant.

Overall, our data shows G-EDF-MP to scale much better than either GSN-EDF or SD in the tested range of up to 64 cores. Nonetheless, G-EDF-MP is obviously not arbitrarily scalable as the DSP eventually becomes a bottleneck, albeit at much higher processor counts than in prior lock-based schedulers.

Next, we discuss whether G-EDF-MP's improved overhead scalability translates into analytical benefits as well.

V. SCHEDULABILITY EXPERIMENTS

Since message passing introduces new sources of delays, just comparing the magnitude of overheads does not provide a complete view of G-EDF-MP's performance. Rather, *overhead-aware* schedulability analysis that accounts for all runtime overheads is necessary. We therefore performed extensive overhead-aware schedulability experiments to assess the impact of overheads on each of the four implementations. Before presenting the main results, we briefly summarize our experimental setup.

A. Experimental Setup

We evaluated the four schedulers (GSN-EDF, GSN-EDF-DI, SD and G-EDF-MP) w.r.t. two criteria: *soft* real-time schedulability, by checking if a task set has *bounded tardiness* [23] after accounting for *average-case* overheads, and *hard* real-time schedulability, by applying a collection of G-EDF schedulability tests [2, 5, 10, 27] after accounting for *worst-case* overheads. Runtime overheads were accounted for using standard techniques as described in Appendix B. For increased breadth, we employed two different task set generation methods; one used in prior LITMUS^{RT} studies [13, 16, 18] and one proposed by Emberson *et al.* [24].

First, we evaluated schedulability as a function of n using task sets similar to those used in previous LITMUS^{RT} studies (see [13, Ch. 4] for a detailed description). In short, each task set is generated by independently drawing n utilizations and

periods. Utilizations were drawn from one of six distributions: three exponential distributions ranging over [0, 1] (with means 0.1, 0.25 and 0.5), and three uniform distributions (ranging over [0.1, 0.2], [0.1, 0.4] and [0.5, 0.9]). Periods were chosen from uniform and log-uniform distributions across several ranges ([1, 1000], [10, 100], [10, 1000], [3, 33], and [200, 1000], in *ms*).

In the second set of experiments using Emberson *et al.*'s task set generation method [24], we assessed the impact of increasing the total utilization for a fixed n. We varied $n \in \{2m, 3m, \ldots, 10m\}$ and considered both uniformly and log-uniformly distributed periods in the range [1, 1000].

Both methods were used to assess the schedulability for each $m \in \{8, 16, 24, 32, 48, 64\}$. For each sampling point, we generated and tested 1,000 task sets.

In total, we evaluated more than 1,000 parameter configurations. Due to space constraints, we highlight here two configurations that reveal a significant analytical impact. The complete set of results is available in Appendix E.

B. Results

Fig. 8(a) shows a result depicting soft real-time schedulability for m = 64 as a function of n. The achieved soft real-time schedulability is shown for each of the four evaluated implementations and also for two hypothetical G-EDF configurations assuming zero overheads on m and m - 1 processors, to provide an upper bound and to show the cost of using a DSP. In Fig. 8(a), GSN-EDF and SD achieve zero schedulability for all task counts due to the pessimism involved in interrupt accounting (recall that GSN-EDF-DI and G-EDF-MP shield real-time tasks from all interrupts, whereas tasks may be delayed by release interrupts in SD and GSN-EDF). Under G-EDF-MP, bounded tardiness is maintained up to n = 448, whereas GSN-EDF-DI starts to decay already at n = 128 and reaches zero schedulability at n = 320. Here, G-EDF-MP can support more than 120 additional tasks, thereby halving the distance to the theoretical upper bound.



Fig. 7: Additional results showing overheads as a function of increasing n for m = 32. Note that each graph uses a different scale due to the wide range of observed overhead magnitudes.

Fig. 8(b) shows hard real-time schedulability as a function of total utilization (for a fixed n). As in Fig. 8(a), interrupt-related pessimism renders GSN-EDF and SD completely unschedulable. About one third of the task sets are schedulable under GSN-EDF-DI until reaching a total utilization of about 20. In contrast, about 70% of the task sets are schedulable under G-EDF-MP until reaching a total utilization of about 24, where the theoretical upper bound starts to decrease as well. Again, G-EDF-MP halves the gap to the theoretical upper bound.

Fig. 8(b) also exhibits G-EDF's inherent hard real-time limitations: even when assuming overhead-free execution, hardly any task set with a total utilization greater than 50% can be claimed schedulable. However, in the context of this work, G-EDF's *algorithmic* limitations are irrelevant because G-EDF-MP's low-overhead design can be easily transferred to other global schedulers (such as global fixed-priority scheduling or *earliest-deadline-until-zero-laxity scheduling* (EDZL) [3], which can offer higher hard real-time schedulability).

Overall, the schedulability experiments confirm that G-EDF-MP's lower overheads translate into significantly higher schedulability at high core counts. At lower core counts (*i.e.*, $m \le 16$) and with fewer tasks, overhead scalability is not yet an issue and G-EDF-MP performs similarly to GSN-EDF-DI.

VI. RELATED WORK

This work reuses two key implementation techniques and is closely related to several prior studies, which we review next.

Implementation techniques. G-EDF-MP employs a DSP to shield real-time tasks from OS overheads, a well-known technique that dates back to the Spring kernel [37]. More recently, the idea of delegating cores to exclusively handle OS duties has also been applied in the context of manycore platforms (*e.g.*, [4, 38]).

G-EDF-MP further employs message passing, which is known to scale well on multicores. In particular, in *multikernel* OSs such as Barrelfish [9] and Quest-V [33], the entire OS design is centered around message passing to avoid scalability bottlenecks



(a) Soft schedulability as a function of n for exponentially distributed task utilizations with mean 0.1 (as in [13, 16, 18]).



(b) Hard schedulability as a function of total utilization based on Emberson *et al.*'s method [24] with n = 5m.

Fig. 8: Comparison of hard and soft schedulability for m = 64. Periods were drawn from a log-uniform distribution with range [1ms,1000ms].

(among other reasons). Message passing is further used in some current commercial RTOSs as well (*e.g.*, ENEA's OSE). However, G-EDF-MP differs significantly from these prior works in that it implements a *global* scheduling policy, wheres OSE, Quest-V, and Barrelfish intentionally avoid global policies whenever possible and implement partitioned schedulers (which scale trivially w.r.t. overheads).

Message passing has also been shown to offer higher throughput than shared locks in the face of heavy contention in recent work studying synchronization choices on multicore platforms [21, 34]. However, throughput depends mostly on averagecase overheads, whereas we designed G-EDF-MP specifically to lower worst-case overheads. Prior to G-EDF-MP, it was not clear that message passing could be a viable—or even superior choice for constructing global real-time schedulers.

Prior studies. Both GSN-EDF and SD have been the subject of prior studies. Most recently, Lelli *et al.* [32] empirically compared SD and stock Linux's SCHED_FIFO scheduler (in various configurations to emulate global, clustered, and partitioned scheduling) on a 48-core platform. In contrast to this study, they did not measure worst-case overheads (and did not apply schedulability analysis), but rather measured average-case overheads and indirectly assessed each scheduler's performance with a statistical evaluation of observed task behavior. Since SD exhibits low average-case overheads, their experimental setup did not reveal the pathological lock contention that we observed.

Several alternative implementations of G-EDF in LITMUS^{RT} were compared in [16]. While several of the investigated alternatives aimed at reducing lock contention, they did not perform significantly better than the stock GSN-EDF plugin, which in hindsight can be attributed to the fact that they did not reduce peak contention (a single lock to serialize scheduling decisions remained in all versions). Dedicated interrupt handling was shown to be very beneficial, which is apparent in Fig. 8, too.

GSN-EDF was previously evaluated as part of several largescale, overhead-aware schedulability studies [7, 13, 18], as was a prior version of G-EDF in LITMUS^{RT} [20]. However, these studies were conducted on platforms significantly smaller than the 64-core machine employed herein, and also considered only a single core count each; scalability limitations hence did not manifest to the extent revealed in this paper. Further, prior studies [7, 13, 18, 20] did not compare LITMUS^{RT} against SD and used statistical outlier filters (in contrast to this study).

Finally, it should be noted that scheduling overheads are just one piece of the predictability puzzle on multicores, as other factors such as contention and interference at all levels of the memory hierarchy must be considered, too (*e.g.*, see [29, 35]).

VII. CONCLUSION

We evaluated the scalability of G-EDF implementations and found that G-EDF-MP, a novel scheduler design based primarily on message passing, scales substantially better than prior lockbased schedulers, thereby lowering scheduling overheads by up to a factor of 36 on 64 cores. While global scheduling is certainly not the right approach for *all* real-time workloads (*e.g.*, CPMD issues remain), this work significantly extends the range of workloads and platforms for which global approaches are viable, thus opening up interesting avenues for future work. In particular, the presented approach is not limited to G-EDF, but can be easily applied to other global scheduling policies as well.

The key limitation of G-EDF-MP is that a single DSP is an inherent bottleneck that inevitably will cause problems beyond 64 cores. Given this limitation, systems with core counts too large to be efficiently supported by a single G-EDF-MP instance should use clustered scheduling instead (whereby the set of all cores is partitioned into disjoint subsets and a global scheduler

is instantiated for each such subset). In fact, this was already noted in a prior study [7], which recommended that global scheduling should be restricted to a "small-to-medium number of cores." Our work does not invalidate this observation, but rather redefines the notion of what constitutes a "medium number of cores." That is, seen in this context, G-EDF-MP simply increases the largest cluster size that can be efficiently supported.

In future work, it would be interesting to develop G-EDF-MP extensions with two or more DSPs, which then, however, would have to synchronize decisions. A substantial algorithmic improvement would be to equip clients with a next pointer to avoid idle time after job completions/suspensions, but choosing next such that analytical benefits are realized is not trivial.

REFERENCES

- [1] "The LITMUS^{RT} project web site," http://www.litmus-rt.org.
- [2] T. Baker and S. Baruah, "An analysis of global EDF schedulability for arbitrary-deadline sporadic task systems," *Real-Time Systems*, vol. 43, no. 1, pp. 3–24, 2009.
- [3] T. Baker, M. Cirinei, and M. Bertogna, "EDZL scheduling analysis," *Real-Time Systems*, vol. 40, no. 3, pp. 264–289, 2008.
- [4] F. Ballesteros, N. Evans, C. Forsyth, G. Guardiola, J. McKie, R. Minnich, and E. Soriano-Salvador, "Nix: a case for a manycore system for cloud computing," *Bell Labs Tech. Journal*, vol. 17, no. 2, pp. 41– 54, 2012.
- [5] S. Baruah, "Techniques for multiprocessor global schedulability analysis," in *RTSS'07*, 2007.
- [6] S. Baruah, J. Gehrke, and C. Plaxton, "Fast scheduling of periodic tasks on multiple resources," in *IPPS*'95, 1995.
- [7] A. Bastoni, B. Brandenburg, and J. Anderson, "An empirical comparison of global, partitioned, and clustered multiprocessor EDF schedulers," in *RTSS'10*, 2010, pp. 14–24.
- [8] —, "Cache-related preemption and migration delays: Empirical approximation and impact on schedulability," in OSPERT'10, 2010.
- [9] A. Baumann, P. Barham, P.-E. Dagand, T. Harris, R. Isaacs, S. Peter, T. Roscoe, A. Schüpbach, and A. Singhania, "The multikernel: a new OS architecture for scalable multicore systems," in SOSP'09, 2009.
- [10] M. Bertogna, M. Cirinei, and G. Lipari, "Improved schedulability analysis of EDF on multiprocessor platforms," in *ECRTS*'05, 2005.
- [11] A. Block, H. Leontyev, B. Brandenburg, and J. Anderson, "A flexible real-time locking protocol for multiprocessors," in *RTCSA'07*, 2007.
- [12] A. Block, "Adaptive multiprocessor real-time systems," Ph.D. dissertation, The University of North Carolina at Chapel Hill, 2008.
- [13] B. Brandenburg, "Scheduling and locking in multiprocessor realtime operating systems," Ph.D. dissertation, The University of North Carolina at Chapel Hill, 2011.
- [14] —, "Improved analysis and evaluation of real-time semaphore protocols for P-FP scheduling," in *RTAS'13*, 2013, pp. 292–302.
- [15] B. Brandenburg and J. Anderson, "Feather-trace: A light-weight event tracing toolkit," in OSPERT'07, 2007.
- [16] —, "On the implementation of global real-time schedulers," in *RTSS'09*, 2009.
- [17] B. Brandenburg, A. Block, J. Calandrino, U. Devi, H. Leontyev, and J. Anderson, "LITMUS^{RT}: A status report," in *RTLWS'07*, 2007.
- [18] B. Brandenburg, J. Calandrino, and J. Anderson, "On the scalability of real-time scheduling algorithms on multicore platforms: A case study," in *RTSS'08*, 2008.
- [19] B. Brandenburg, H. Leontyev, and J. Anderson, "An overview of interrupt accounting techniques for multiprocessor real-time systems," *Journal of Systems Architecture: Embedded Software Design*, vol. 57, no. 6, pp. 638–654, 2011.
- [20] J. Calandrino, H. Leontyev, A. Block, U. Devi, and J. Anderson, "LITMUS^{RT}: A testbed for empirically comparing real-time multiprocessor schedulers," in *RTSS'06*, 2006.
- [21] T. David, R. Guerraoui, and V. Trigonakis, "Everything you always wanted to know about synchronization but were afraid to ask," in *SOSP'13*, 2013.
- [22] R. Davis and A. Burns, "A survey of hard real-time scheduling for multiprocessor systems," ACM Comput. Surv., vol. 43, no. 4, pp. 35:1–35:44, 2011.

- [23] U. Devi and J. Anderson, "Tardiness bounds under global EDF scheduling on a multiprocessor," *Real-Time Systems*, vol. 38, no. 2, pp. 133–189, 2008.
- [24] P. Emberson, R. Stafford, and R. Davis, "Techniques for the synthesis of multiprocessor tasksets," WATERS'10, 2010.
- [25] J. Erickson and J. Anderson, "Fair lateness scheduling: Reducing maximum lateness in G-EDF-like scheduling," in *ECRTS*'12, 2012.
- [26] D. Faggioli, F. Checconi, M. Trimarchi, and C. Scordino, "An EDF scheduling class for the Linux kernel," in *RTLWS'09*, 2009.
- [27] J. Goossens, S. Funk, and S. Baruah, "Priority-driven scheduling of periodic task systems on multiprocessors," *Real-Time Systems*, vol. 25, no. 2-3, pp. 187–205, 2003.
- [28] A. Gujarati, F. Cerqueira, and B. Brandenburg, "Schedulability analysis of the Linux push and pull scheduler with arbitrary processor affinities," in *ECRTS'13*, 2013.
- [29] Y. Heechul, G. Yao, R. Pellizzoni, M. Caccamo, and L. Sha, "Mem-Guard: Memory bandwidth reservation system for efficient performance isolation in multi-core platforms," in *RTAS'13*, 2013.
- [30] V. Legout, M. Jan, and L. Pautet, "A scheduling algorithm to reduce the static energy consumption of multiprocessor real-time systems," in *RTNS'13*, 2013.
- [31] J. Lelli, G. Lipari, D. Faggioli, and T. Cucinotta, "An efficient and scalable implementation of global EDF in Linux," in OSPERT'11, 2011.
- [32] J. Lelli, D. Faggioli, T. Cucinotta, and G. Lipari, "An experimental comparison of different real-time schedulers on multicore systems," *J. Syst. Softw.*, vol. 85, no. 10, pp. 2405–2416, 2012.
- [33] Y. Li, M. Danish, and R. West, "Quest-V: A virtualized multikernel for high-confidence systems," 2011, arXiv:1112.5136.
- [34] J.-P. Lozi, F. David, G. Thomas, J. Lawall, and G. Muller, "Remote core locking: Migrating critical-section execution to improve the performance of multithreaded applications," in *ATC'12*, 2012.
- [35] R. Mancuso, R. Dudko, E. Betti, M. Cesati, M. Caccamo, and R. Pellizzoni, "Real-time cache management framework for multicore architectures," in *RTAS'13*, 2013.
- [36] P. Regnier, G. Lima, E. Massa, G. Levin, and S. Brandt, "RUN: Optimal multiprocessor real-time scheduling via reduction to uniprocessor," in *RTSS'11*, 2011.
- [37] J. Stankovic and K. Ramamritham, "The Spring kernel: a new paradigm for real-time operating systems," *SIGOPS Oper. Syst. Rev.*, vol. 23, no. 3, pp. 54–71, Jul. 1989.
- [38] D. Wentzlaff and A. Agarwal, "Factored operating systems (fos): the case for a scalable operating system for multicores," ACM SIGOPS Operating Systems Review, vol. 43, no. 2, pp. 76–85, 2009.

APPENDIX

A. Overhead Experiments

Overheads under LITMUS^{RT} are commonly measured using workloads consisting of synthetic periodic tasks with randomly generated task parameters (*i.e.*, execution costs and periods) [13, 16, 18]. These benchmark tasks carry out floating point arithmetic on large arrays to simulate CPU-bound signal processing tasks. We adopted this approach in this study as well, but needed to port the tasks to SD to enable a fair comparison.

LITMUS^{RT} provides custom system calls to support periodic job releases and synchronous task set releases (which means that all tasks release their first job at exactly the same time), which we emulated under SD using standard Linux APIs. Periodic execution can be trivially accomplished under Linux with clock_nanosleep(). Synchronous task set releases were coordinated with a shared memory segment and *pthreads* barriers, as illustrated in Fig. 9. The benchmark workload hence behaves virtually identically under SD and LITMUS^{RT}.

B. Overhead Accounting

Runtime overheads can be accounted for in schedulability tests that assume overhead-free execution (*i.e.*, that consider

```
barrier_wait();
                      // await synchronous release
1
2
   next_rel = start; // the shared release time
3
   while (!timeout()) {
       clock_nanosleep (CLOCK_MONOTONIC,
4
            TIMER_ABSTIME,
                               &next_rel, NULL);
       // Do some math and dirty the cache ...
5
6
       next_rel = timespec_add(next_rel,
            period_timespec);
7
   }
```

Fig. 9: Skeleton of the periodic tasks used to benchmark the SD scheduler.

overheads to be negligible) by inflating task parameters so that a positive overhead-free schedulability test of the transformed task set implies the timeliness of the original task set in the presence of overheads (*e.g.*, context switch costs can be charged to the preempting job as execution costs, *etc.*). The required transformations are conceptually simple, but somewhat tedious and beyond the scope of this paper; the interested reader may find a detailed description of the standard techniques applicable to GSN-EDF, GSN-EDF-DI, and SD in [13, Ch. 3].

However, G-EDF-MP requires additional accounting for the communication that it introduces. For the sake of completeness, we briefly document how we extended the accounting methods provided in [13] to reflect the added delays. We assume familiarity with the basic approach on behalf of the reader.

From an overhead accounting perspective, G-EDF-MP works similarly to GSN-EDF-DI. Thus, the accounting methods for scheduling (Δ^{sch} , Δ^{ipi}) and interrupt overheads (Δ^{rel} , Δ^{ev}) do not differ. However, every time a preemption occurs, the preempted lower-priority job is able to resume execution only after a full round-trip communication with the DSP. That is, (i) the client becomes idle, (ii) signals the completion, (iii) the DSP links a new task to the client and sends back an IPI. This situation is illustrated in Fig. 5, where the client processor becomes idle while task T_1 is pending.

Compared to the overheads delaying a preempted job under GSN-EDF or GSN-EDF-DI, this introduces additional delays due to the client request latency (Δ^{req}), the DSP handler (Δ^{dsp}), an additional IPI delay (Δ^{ipi}), and an additional scheduler invocation (Δ^{sch}) and context switch (Δ^{cxs}). This results in the following rule for inflating parameters.

For a sporadic task T_i with period p_i and execution cost e_i , let e'_i and p'_i denote the transformed parameters. Then

 $p_i' = p_i - \Delta^{ev}$

and

$$e'_{i} = \frac{e_{i} + 3 \cdot (\Delta^{sch} + \Delta^{cxs}) + \Delta^{cpd}}{1 - U_{0}^{tck}} + 2 \cdot C^{pre} + \Delta^{req} + \Delta^{dsp} + 2 \cdot \Delta^{ipi} + \Delta^{rel}$$

where $C^{pre} = \frac{\Delta^{tck} + \Delta^{ev} \cdot U^{tck}}{1 - U^{tck}}$ denotes the interrupt-related cost per preemption [13, Sec. 3.4.5], and where $U^{tck} = \frac{\Delta^{tck}}{1000\mu s}$ denotes the utilization of the tick interrupt.

C. All Overhead Graphs

This section contains the complete set of overhead measurements that we observed in our experiments. As mentioned in Sec. IV, we collected the overheads under four different implementations of G-EDF: G-EDF-MP, SD, GSN-EDF *without* dedicated interrupt handling, and GSN-EDF *with* dedicated interrupt handling (GSN-EDF-DI).

Each scheduler was executed with different processor counts $m \in \{8, 16, 24, 32, 48, 64\}$. For each m and for each $n \in \{m, 2m, \ldots, 10m\}$, we generated five task sets with n synthetic tasks using Emberson *et al.*'s task set generator [24]. Task sets were generated with a target utilization chosen uniformly at random from [0.4m, 0.8m] and periods drawn from a log-uniform distribution in the range [10ms, 100ms].

We measured the following sources of runtime overheads:

- 1) scheduling overhead;
- 2) release overhead;
- 3) release latency;
- 4) tick overhead;
- 5) context-switch overhead;
- 6) IPI latency; and specific to G-EDF-MP,
- 7) client request latency; and
- 8) DSP handler overhead.

Since scheduling overhead in LITMUS^{RT} is measured as two parts (before and after a context switch), we report each part individually (denoted **1a** and **1b**).

For each overhead and each considered processor count, we show graphs with the maximum and average observed values. No statistical outlier filter was applied. The graphs are organized as follows.

- Figs. 10–25 show overheads for m = 8.
- Figs. 26–41 show overheads for m = 16.
- Figs. 42–57 show overheads for m = 24.
- Figs. 58–73 show overheads for m = 32.
- Figs. 74–89 show overheads for m = 48.
- Figs. 90–105 show overheads for m = 64.



Fig. 10: Maximum observed scheduling overhead (a) for 8 processors as a function of the number of tasks.



Fig. 11: Average observed scheduling overhead (a) for 8 processors as a function of the number of tasks.



Fig. 12: Maximum observed scheduling overhead (b) for 8 processors as a function of the number of tasks.



Fig. 13: Average observed scheduling overhead (b) for 8 processors as a function of the number of tasks.



Fig. 14: Maximum observed release overhead for 8 processors as a function of the number of tasks.



Fig. 15: Average observed release overhead for 8 processors as a function of the number of tasks.



Fig. 16: Maximum observed release latency for 8 processors as a function of the number of tasks.



Fig. 17: Average observed release latency for 8 processors as a function of the number of tasks.



Fig. 18: Maximum observed tick overhead for 8 processors as a function of the number of tasks.



Fig. 19: Average observed tick overhead for 8 processors as a function of the number of tasks.



Fig. 20: Maximum observed context-switching overhead for 8 processors as a function of the number of tasks.



Fig. 21: Average observed context-switching overhead for 8 processors as a function of the number of tasks.



Fig. 22: Maximum observed IPI latency for 8 processors as a function of the number of tasks.



Fig. 23: Average observed IPI latency for 8 processors as a function of the number of tasks.



Fig. 24: Maximum and average observed client request latency for 8 processors as a function of the number of tasks.



Fig. 25: Maximum and average observed DSP handler overhead for 8 processors as a function of the number of tasks.



Fig. 26: Maximum observed scheduling overhead (a) for 16 processors as a function of the number of tasks.



Fig. 27: Average observed scheduling overhead (a) for 16 processors as a function of the number of tasks.



Fig. 28: Maximum observed scheduling overhead (b) for 16 processors as a function of the number of tasks.



Fig. 29: Average observed scheduling overhead (b) for 16 processors as a function of the number of tasks.



Fig. 30: Maximum observed release overhead for 16 processors as a function of the number of tasks.



Fig. 31: Average observed release overhead for 16 processors as a function of the number of tasks.



Fig. 32: Maximum observed release latency for 16 processors as a function of the number of tasks.



Fig. 33: Average observed release latency for 16 processors as a function of the number of tasks.



Fig. 34: Maximum observed tick overhead for 16 processors as a function of the number of tasks.



Fig. 35: Average observed tick overhead for 16 processors as a function of the number of tasks.



Fig. 36: Maximum observed context-switching overhead for 16 processors as a function of the number of tasks.



Fig. 37: Average observed context-switching overhead for 16 processors as a function of the number of tasks.



Fig. 38: Maximum observed IPI latency for 16 processors as a function of the number of tasks.



Fig. 39: Average observed IPI latency for 16 processors as a function of the number of tasks.



Fig. 40: Maximum and average observed client request latency for 16 processors as a function of the number of tasks.



Fig. 41: Maximum and average observed DSP handler overhead for 16 processors as a function of the number of tasks.



Fig. 42: Maximum observed scheduling overhead (a) for 24 processors as a function of the number of tasks.



Fig. 43: Average observed scheduling overhead (a) for 24 processors as a function of the number of tasks.



Fig. 44: Maximum observed scheduling overhead (b) for 24 processors as a function of the number of tasks.



Fig. 45: Average observed scheduling overhead (b) for 24 processors as a function of the number of tasks.



Fig. 46: Maximum observed release overhead for 24 processors as a function of the number of tasks.



Fig. 47: Average observed release overhead for 24 processors as a function of the number of tasks.



Fig. 48: Maximum observed release latency for 24 processors as a function of the number of tasks.



Fig. 49: Average observed release latency for 24 processors as a function of the number of tasks.



Fig. 50: Maximum observed tick overhead for 24 processors as a function of the number of tasks.



Fig. 51: Average observed tick overhead for 24 processors as a function of the number of tasks.



Fig. 52: Maximum observed context-switching overhead for 24 processors as a function of the number of tasks.



Fig. 53: Average observed context-switching overhead for 24 processors as a function of the number of tasks.



Fig. 54: Maximum observed IPI latency for 24 processors as a function of the number of tasks.



Fig. 55: Average observed IPI latency for 24 processors as a function of the number of tasks.



Fig. 56: Maximum and average observed client request latency for 24 processors as a function of the number of tasks.



Fig. 57: Maximum and average observed DSP handler overhead for 24 processors as a function of the number of tasks.



Fig. 58: Maximum observed scheduling overhead (a) for 32 processors as a function of the number of tasks.



Fig. 59: Average observed scheduling overhead (a) for 32 processors as a function of the number of tasks.



Fig. 60: Maximum observed scheduling overhead (b) for 32 processors as a function of the number of tasks.



Fig. 61: Average observed scheduling overhead (b) for 32 processors as a function of the number of tasks.



Fig. 62: Maximum observed release overhead for 32 processors as a function of the number of tasks.



Fig. 63: Average observed release overhead for 32 processors as a function of the number of tasks.



Fig. 64: Maximum observed release latency for 32 processors as a function of the number of tasks.



Fig. 65: Average observed release latency for 32 processors as a function of the number of tasks.



Fig. 66: Maximum observed tick overhead for 32 processors as a function of the number of tasks.



Fig. 67: Average observed tick overhead for 32 processors as a function of the number of tasks.



Fig. 68: Maximum observed context-switching overhead for 32 processors as a function of the number of tasks.



Fig. 69: Average observed context-switching overhead for 32 processors as a function of the number of tasks.



Fig. 70: Maximum observed IPI latency for 32 processors as a function of the number of tasks.



Fig. 71: Average observed IPI latency for 32 processors as a function of the number of tasks.



Fig. 72: Maximum and average observed client request latency for 32 processors as a function of the number of tasks.



Fig. 73: Maximum and average observed DSP handler overhead for 32 processors as a function of the number of tasks.



Fig. 74: Maximum observed scheduling overhead (a) for 48 processors as a function of the number of tasks.



Fig. 75: Average observed scheduling overhead (a) for 48 processors as a function of the number of tasks.



Fig. 76: Maximum observed scheduling overhead (b) for 48 processors as a function of the number of tasks.



Fig. 77: Average observed scheduling overhead (b) for 48 processors as a function of the number of tasks.



Fig. 78: Maximum observed release overhead for 48 processors as a function of the number of tasks.



Fig. 79: Average observed release overhead for 48 processors as a function of the number of tasks.



Fig. 80: Maximum observed release latency for 48 processors as a function of the number of tasks.



Fig. 81: Average observed release latency for 48 processors as a function of the number of tasks.



Fig. 82: Maximum observed tick overhead for 48 processors as a function of the number of tasks.



Fig. 83: Average observed tick overhead for 48 processors as a function of the number of tasks.



Fig. 84: Maximum observed context-switching overhead for 48 processors as a function of the number of tasks.



Fig. 85: Average observed context-switching overhead for 48 processors as a function of the number of tasks.



Fig. 86: Maximum observed IPI latency for 48 processors as a function of the number of tasks.



Fig. 87: Average observed IPI latency for 48 processors as a function of the number of tasks.



Fig. 88: Maximum and average observed client request latency for 48 processors as a function of the number of tasks.



Fig. 89: Maximum and average observed DSP handler overhead for 48 processors as a function of the number of tasks.



Fig. 90: Maximum observed scheduling overhead (a) for 64 processors as a function of the number of tasks.



Fig. 91: Average observed scheduling overhead (a) for 64 processors as a function of the number of tasks.



Fig. 92: Maximum observed scheduling overhead (b) for 64 processors as a function of the number of tasks.



Fig. 93: Average observed scheduling overhead (b) for 64 processors as a function of the number of tasks.



Fig. 94: Maximum observed release overhead for 64 processors as a function of the number of tasks.



Fig. 95: Average observed release overhead for 64 processors as a function of the number of tasks.



Fig. 96: Maximum observed release latency for 64 processors as a function of the number of tasks.



Fig. 97: Average observed release latency for 64 processors as a function of the number of tasks.



Fig. 98: Maximum observed tick overhead for 64 processors as a function of the number of tasks.



Fig. 99: Average observed tick overhead for 64 processors as a function of the number of tasks.



Fig. 100: Maximum observed context-switching overhead for 64 processors as a function of the number of tasks.



Fig. 101: Average observed context-switching overhead for 64 processors as a function of the number of tasks.



Fig. 102: Maximum observed IPI latency for 64 processors as a function of the number of tasks.



Fig. 103: Average observed IPI latency for 64 processors as a function of the number of tasks.



Fig. 104: Maximum and average observed client request latency for 64 processors as a function of the number of tasks.



Fig. 105: Maximum and average observed DSP handler overhead for 64 processors as a function of the number of tasks.

D. All Overhead Histograms

This section contains the complete set of histograms for the observed overhead measurements, considering the four different implementations of G-EDF that were evaluated: G-EDF-MP, SD, GSN-EDF *without* dedicated interrupt handling, and GSN-EDF *with* dedicated interrupt handling (GSN-EDF-DI).

The histograms depict the same data used for the overhead graphs, whose experimental setup can be found in Appendix C.

For each overhead and each considered processor count, we present the respective normalized histograms. The values of minimum, maximum, average, median and standard deviation are also included. No statistical outlier filter was applied. The histograms are organized as follows.

- Figs. 106–405 show results for m = 8.
 - Figs. 106–135 show results for n = 1m = 8.
 - Figs. 136–165 show results for n = 2m = 16.
 - Figs. 166–195 show results for n = 3m = 24.
 - Figs. 196–225 show results for n = 4m = 32.
 - Figs. 226–255 show results for n = 5m = 40.
 - Figs. 256–285 show results for n = 6m = 48.
 - Figs. 286–315 show results for n = 7m = 56.
 - Figs. 316–345 show results for n = 8m = 64.
 - Figs. 346–375 show results for n = 9m = 72.
 - Figs. 376–405 show results for n = 10m = 80.
- Figs. 406–705 show results for m = 16.
 - Figs. 406–435 show results for n = 1m = 16.
 - Figs. 436–465 show results for n = 2m = 32.
 - Figs. 466–495 show results for n = 3m = 48.
 - Figs. 496–525 show results for n = 4m = 64.
 - Figs. 526–555 show results for n = 5m = 80.
 - Figs. 556–585 show results for n = 6m = 96.
 - Figs. 586–615 show results for n = 7m = 112.
 - Figs. 616–645 show results for n = 8m = 128.
 - Figs. 646–675 show results for n = 9m = 144.
 - Figs. 676–705 show results for n = 10m = 160.
- Figs. 706–1005 show results for m = 24.
 - Figs. 706–735 show results for n = 1m = 24.
 - Figs. 736–765 show results for n = 2m = 48.
 - Figs. 766–795 show results for n = 3m = 72.
 - Figs. 796–825 show results for n = 4m = 96.
 - Figs. 826–855 show results for n = 5m = 120.
 - Figs. 856–885 show results for n = 6m = 144.
 - Figs. 886–915 show results for n = 7m = 168.
 - Figs. 916–945 show results for n = 8m = 192.
 - Figs. 946–975 show results for n = 9m = 216.
 - Figs. 976–1005 show results for n = 10m = 240.
- Figs. 1006–1305 show results for m = 32.
 - Figs. 1006–1035 show results for n = 1m = 32.
 - Figs. 1036–1065 show results for n = 2m = 64.
 - Figs. 1066–1095 show results for n = 3m = 96.
 - Figs. 1096–1125 show results for n = 4m = 128.
 - Figs. 1126–1155 show results for n = 5m = 160.
 - Figs. 1156–1185 show results for n = 6m = 192.
 - Figs. 1186–1215 show results for n = 7m = 224.
 - Figs. 1216–1245 show results for n = 8m = 256.

- Figs. 1246–1275 show results for n = 9m = 288.
- Figs. 1276–1305 show results for n = 10m = 320.
- Figs. 1306–1605 show results for m = 48.
 - Figs. 1306–1335 show results for n = 1m = 48.
 - Figs. 1336–1365 show results for n = 2m = 96.
 - Figs. 1366–1395 show results for n = 3m = 144.
 - Figs. 1396–1425 show results for n = 4m = 192.
 - Figs. 1426–1455 show results for n = 5m = 240.
 - Figs. 1456–1485 show results for n = 6m = 288.
 - Figs. 1486–1515 show results for n = 7m = 336.
 - Figs. 1516–1545 show results for n = 8m = 384.
 - Figs. 1546–1575 show results for n = 9m = 432.
 - Figs. 1576–1605 show results for n = 10m = 480.
- Figs. 1606–1905 show results for m = 64.
 - Figs. 1606–1635 show results for n = 1m = 64.
 - Figs. 1636–1665 show results for n = 2m = 128.
 - Figs. 1666–1695 show results for n = 3m = 192.
 - Figs. 1696–1725 show results for n = 4m = 256.
 - Figs. 1726–1755 show results for n = 5m = 320.
 - Figs. 1756–1785 show results for n = 6m = 384.
 - Figs. 1786–1815 show results for n = 7m = 448.
 - Figs. 1816–1845 show results for n = 8m = 512.
 - Figs. 1846–1875 show results for n = 9m = 576.
 - Figs. 1876–1905 show results for n = 10m = 640.



Fig. 106: Histogram of scheduling overhead (a) under GSN-EDF, with m = 8 and n = 8.



Fig. 107: Histogram of scheduling overhead (b) under GSN-EDF, with m = 8 and n = 8.



Fig. 108: Histogram of release overhead under GSN-EDF, with m = 8and n = 8.



Fig. 109: Histogram of release latency under GSN-EDF, with m = 8and n = 8.

samples: total=561168 fraction of samples 0.1 0.01 0.001 0.0001 0.05 0.25 0.45 0.65 0.85 1.05 1.25 1.45

Fig. 110: Histogram of tick overhead under GSN-EDF, with m = 8 and n = 8.



overhead in microseconds (bin size = 0.30us) Fig. 111: Histogram of context-switching overhead under GSN-EDF, with m = 8 and n = 8.









overhead in microseconds (bin size = 0.10us)





Fig. 114: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 8 and n = 8.



Fig. 115: Histogram of release overhead under GSN-EDF-DI, with m = 8 and n = 8.



Fig. 116: Histogram of release latency under GSN-EDF-DI, with m = 8 and n = 8.



Fig. 117: Histogram of tick overhead under GSN-EDF-DI, with m = 8 and n = 8.



Fig. 118: Histogram of context-switching overhead under GSN-EDF-DI, with m = 8 and n = 8.





Fig. 119: Histogram of IPI latency under GSN-EDF-DI, with m = 8 and n = 8.











Fig. 122: Histogram of release overhead under G-EDF-MP, with m = 8 and n = 8.



Fig. 123: Histogram of release latency under G-EDF-MP, with m = 8 and n = 8.



Fig. 124: Histogram of tick overhead under G-EDF-MP, with m = 8 and n = 8.



Fig. 125: Histogram of context-switching overhead under G-EDF-MP, with m = 8 and n = 8.



Fig. 126: Histogram of IPI latency under G-EDF-MP, with m = 8 and n = 8.



Fig. 127: Histogram of client request latency under G-EDF-MP, with m = 8 and n = 8.







Fig. 129: Histogram of scheduling overhead (a) under SD, with m = 8 and n = 8.



Fig. 130: Histogram of scheduling overhead (b) under SD, with m = 8 and n = 8.



Fig. 131: Histogram of release overhead under SD, with m = 8 and n = 8.



Fig. 132: Histogram of release latency under SD, with m = 8 and n = 8.



Fig. 133: Histogram of tick overhead under SD, with m = 8 and n = 8.

min=0.28us max=3.97us avg=1.36us median=1.36us stdev=0.44us amples: total=41726 0.01 0.001 0.001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 overhead in microseconds (bin size = 0.30us)

Fig. 134: Histogram of context-switching overhead under SD, with m = 8 and n = 8.



Fig. 135: Histogram of IPI latency under SD, with m = 8 and n = 8.



Fig. 136: Histogram of scheduling overhead (a) under GSN-EDF, with m=8 and n=16.



Fig. 137: Histogram of scheduling overhead (b) under GSN-EDF, with m = 8 and n = 16.



Fig. 138: Histogram of release overhead under GSN-EDF, with m=8 and n=16.



Fig. 139: Histogram of release latency under GSN-EDF, with m = 8 and n = 16.



Fig. 140: Histogram of tick overhead under GSN-EDF, with m = 8 and n = 16.



Fig. 141: Histogram of context-switching overhead under GSN-EDF, with m=8 and n=16.



Fig. 142: Histogram of IPI latency under GSN-EDF, with m = 8 and n = 16.











Fig. 145: Histogram of release overhead under GSN-EDF-DI, with m=8 and n=16.



Fig. 146: Histogram of release latency under GSN-EDF-DI, with m=8and n = 16.



Fig. 147: Histogram of tick overhead under GSN-EDF-DI, with m = 8and n = 16.



Fig. 148: Histogram of context-switching overhead under GSN-EDF-DI, with m = 8 and n = 16.



Fig. 149: Histogram of IPI latency under GSN-EDF-DI, with m = 8and n = 16.

min=0.28us max=7.78us avg=1.08us median=1.02us stdev=0.46us samples: total=109785 fraction of samples 0.1 0.01 0.001 0.0001 0.25 1.25 2.25 3.25 4.25 5.25 6.25 7.25



min=0.07us max=1.06us avg=0.09us median=0.07us stdev=0.05us samples: total=109795 fraction of samples 0.1 0.01 0.001 0.0001 0.05 0.25 0.45 0.65 0.85 1.05 overhead in microseconds (bin size = 0.10us)







min=0.31us max=4.49us avg=0.81us median=0.76us stdev=0.27us samples: total=50297 fraction of samples 0.1 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 4.35 overhead in microseconds (bin size = 0.30us)





Fig. 154: Histogram of tick overhead under G-EDF-MP, with m = 8and n = 16.



Fig. 155: Histogram of context-switching overhead under G-EDF-MP, with m = 8 and n = 16.



Fig. 156: Histogram of IPI latency under G-EDF-MP, with m = 8 and n = 16.



Fig. 157: Histogram of client request latency under G-EDF-MP, with m = 8 and n = 16.

min=0.07us max=7.89us avg=1.25us median=1.13us stdev=0.63us samples: total=98902 fraction of samples 0.1 0.01 0.001 0.0001 0.25 1.25 2.25 3.25 4.25 5.25 6.25 7.25











min=0.28us max=13.92us avg=1.76us median=1.47us stdev=0.97us samples: total=99198 fraction of samples 0. 0.01 0.00 0.0001 0.40 2.00 3.60 5.20 6.80 8.40 10.00 11.60 13.20 overhead in microseconds (bin size = 0.80us)





Fig. 162: Histogram of release latency under SD, with m = 8 and n = 16.



Fig. 163: Histogram of tick overhead under SD, with m = 8 and n = 16.



Fig. 164: Histogram of context-switching overhead under SD, with m = 8 and n = 16.



Fig. 165: Histogram of IPI latency under SD, with m = 8 and n = 16.

min=0.25us max=15.18us avg=1.88us median=1.36us stdev=1.38us samples: total=152162 fraction of samples 0.1 0.01 0.00 0.0001 0.45 2.25 4.05 5.85 7.65 9.45 11.25 13.05 14.85

overhead in microseconds (bin size = 0.90us) Fig. 166: Histogram of scheduling overhead (a) under GSN-EDF, with m = 8 and n = 24.



Fig. 167: Histogram of scheduling overhead (b) under GSN-EDF, with m = 8 and n = 24.





min=0.29us max=6.95us avg=0.77us median=0.71us stdev=0.27us samples: total=65764 fraction of samples 0.1 0.01 0.001 0.0001 0.20 1.00 1.80 2.60 3.40 4.20 5.00 5.80 6.60 overhead in microseconds (bin size = 0.40us)









Fig. 171: Histogram of context-switching overhead under GSN-EDF, with m = 8 and n = 24.



Fig. 172: Histogram of IPI latency under GSN-EDF, with m = 8 and n = 24.



Fig. 173: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m=8 and n=24.

min=0.07us max=1.07us avg=0.08us median=0.07us stdev=0.05us



Fig. 174: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 8 and n = 24.



Fig. 175: Histogram of release overhead under GSN-EDF-DI, with m = 8 and n = 24.





min=0.14us max=1.51us avg=0.24us median=0.22us stdev=0.08us samples: total=498378 fraction of samples 0.1 0.01 0.001 0.0001 0.05 0.25 0.45 0.65 0.85 1.05 1.25 1.45 overhead in microseconds (bin size = 0.10us)





Fig. 178: Histogram of context-switching overhead under GSN-EDF-DI, with m = 8 and n = 24.



Fig. 179: Histogram of IPI latency under GSN-EDF-DI, with m = 8and n = 24.



Fig. 180: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 8 and n = 24.



Fig. 181: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 8 and n = 24.

min=0.46us max=16.50us avg=2.77us median=2.17us stdev=1.97us samples: total=65761 fraction of samples 0. 0.01 0.00 0.0001 0.50 2.50 4.50 6.50 8.50 10.50 12.50 14.50 16.50 overhead in microseconds (bin size = 1.00us)

Fig. 182: Histogram of release overhead under G-EDF-MP, with m = 8and n = 24.



Fig. 183: Histogram of release latency under G-EDF-MP, with m = 8and n = 24.














Fig. 187: Histogram of client request latency under G-EDF-MP, with m = 8 and n = 24.



Fig. 188: Histogram of DSP handler overhead under G-EDF-MP, with m = 8 and n = 24.



Fig. 189: Histogram of scheduling overhead (a) under SD, with m = 8and n = 24.

min=0.06us max=17.93us avg=0.83us median=0.07us stdev=1.38us samples: total=187287 fraction of samples 0.1 0.01 0.00 0.0001 0.50 2.50 4.50 6.50 8.50 10.50 12.50 14.50 16.50 overhead in microseconds (bin size = 1.00us)

Fig. 190: Histogram of scheduling overhead (b) under SD, with m = 8and n = 24.













37



Fig. 194: Histogram of context-switching overhead under SD, with m = 8 and n = 24.



Fig. 195: Histogram of IPI latency under SD, with m = 8 and n = 24.



Fig. 196: Histogram of scheduling overhead (a) under GSN-EDF, with m = 8 and n = 32.



Fig. 197: Histogram of scheduling overhead (b) under GSN-EDF, with m = 8 and n = 32.

min=0.37us max=27.80us avg=3.62us median=2.62us stdev=2.77us samples: total=7066 fraction of samples 0.1 0.01 0.001 0.0001 1.00 5.00 9.00 13.00 17.00 21.00 25.00 overhead in microseconds (bin size = 2.00us)





Fig. 199: Histogram of release latency under GSN-EDF, with m = 8









and n = 32.







Fig. 202: Histogram of IPI latency under GSN-EDF, with m = 8 and n = 32.



Fig. 203: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 8 and n = 32.



Fig. 204: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 8 and n = 32.



Fig. 205: Histogram of release overhead under GSN-EDF-DI, with m = 8 and n = 32.



Fig. 206: Histogram of release latency under GSN-EDF-DI, with m=8 and n=32.



Fig. 207: Histogram of tick overhead under GSN-EDF-DI, with m = 8 and n = 32.





min=0.79us max=5.85us avg=1.32us median=1.29us stdev=0.24us samples: ^Itotal=166850 raction of samples 0.1 0.01 0.001 0.0001 0.20 1.00 1.80 2.60 3.40 4.20 5.00 5.80 overhead in microseconds (bin size = 0.40us)



39



Fig. 210: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 8 and n = 32.



Fig. 211: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 8 and n = 32.



Fig. 212: Histogram of release overhead under G-EDF-MP, with m = 8and n = 32.



Fig. 213: Histogram of release latency under G-EDF-MP, with m = 8and n = 32.

samples: total=597210 fraction of samples 0.1 0.01 0.001 0.0001 0.05 0.25 0.45 0.65 0.85 1.05 1.25 1.45 overhead in microseconds (bin size = 0.10us)

Fig. 214: Histogram of tick overhead under G-EDF-MP, with m = 8and n = 32.













min=0.13us max=1.49us avg=0.22us median=0.18us stdev=0.08us



Fig. 218: Histogram of DSP handler overhead under G-EDF-MP, with m=8 and n=32.



Fig. 219: Histogram of scheduling overhead (a) under SD, with m = 8 and n = 32.



Fig. 220: Histogram of scheduling overhead (b) under SD, with m = 8 and n = 32.



Fig. 221: Histogram of release overhead under SD, with m = 8 and n = 32.

min=0.31us max=18.63us avg=1.17us median=0.74us stdev=1.28us 0.01 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.00 0.001 0.00 0.001 0.00 0.001 0.00 0.001 0.00 0.001 0.00 0.001 0.00 0.001 0.00 0.001 0.00 0.001 0.00 0.001 0.00 0.001 0.00 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.0010.001

Fig. 222: Histogram of release latency under SD, with m = 8 and n = 32.



Fig. 223: Histogram of tick overhead under SD, with m = 8 and n = 32.











Fig. 226: Histogram of scheduling overhead (a) under GSN-EDF, with m = 8 and n = 40.



Fig. 227: Histogram of scheduling overhead (b) under GSN-EDF, with m = 8 and n = 40.



Fig. 228: Histogram of release overhead under GSN-EDF, with m = 8and n = 40.



Fig. 229: Histogram of release latency under GSN-EDF, with m = 8and n = 40.

min=0.14us max=1.85us avg=0.24us median=0.23us stdev=0.08us 1 samples: total=423830 fraction of samples 0.1 0.01 0.001 0.0001 0.10 0.30 0.50 0.70 0.90 1.10 1.30 1.50 1.70 1.90 overhead in microseconds (bin size = 0.20us)

















Fig. 234: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 8 and n = 40.



Fig. 235: Histogram of release overhead under GSN-EDF-DI, with m = 8 and n = 40.



Fig. 236: Histogram of release latency under GSN-EDF-DI, with m = 8and n = 40.



Fig. 237: Histogram of tick overhead under GSN-EDF-DI, with m = 8and n = 40.



Fig. 238: Histogram of context-switching overhead under GSN-EDF-DI, with m = 8 and n = 40.



Fig. 239: Histogram of IPI latency under GSN-EDF-DI, with m=8and n = 40.





min=0.07us max=1.22us avg=0.08us median=0.07us stdev=0.04us samples: total=258287 fraction of samples 0.1 0.01 0.001 0.0001 0.05 0.25 0.45 0.65 0.85 1.05 1.25 overhead in microseconds (bin size = 0.10us)





Fig. 242: Histogram of release overhead under G-EDF-MP, with m=8 and n=40.



Fig. 243: Histogram of release latency under G-EDF-MP, with m = 8 and n = 40.



Fig. 244: Histogram of tick overhead under G-EDF-MP, with m = 8 and n = 40.



Fig. 245: Histogram of context-switching overhead under G-EDF-MP, with m=8 and n=40.



Fig. 246: Histogram of IPI latency under G-EDF-MP, with m = 8 and n = 40.



Fig. 247: Histogram of client request latency under G-EDF-MP, with m = 8 and n = 40.











Fig. 250: Histogram of scheduling overhead (b) under SD, with m = 8and n = 40.



Fig. 251: Histogram of release overhead under SD, with m = 8 and n = 40.



Fig. 252: Histogram of release latency under SD, with m = 8 and n = 40.



Fig. 253: Histogram of tick overhead under SD, with m = 8 and n = 40.



Fig. 254: Histogram of context-switching overhead under SD, with m = 8 and n = 40.



Fig. 255: Histogram of IPI latency under SD, with m = 8 and n = 40.





min=0.07us max=1.18us avg=0.08us median=0.07us stdev=0.05us samples: total=314643 raction of samples 0.1 0.01 0.001 0 0001 0.05 0.25 0.45 0.65 0.85 1.05 overhead in microseconds (bin size = 0.10us)





Fig. 258: Histogram of release overhead under GSN-EDF, with m=8 and n=48.



Fig. 259: Histogram of release latency under GSN-EDF, with m = 8 and n = 48.



Fig. 260: Histogram of tick overhead under GSN-EDF, with m = 8 and n = 48.



Fig. 261: Histogram of context-switching overhead under GSN-EDF, with m = 8 and n = 48.



Fig. 262: Histogram of IPI latency under GSN-EDF, with m = 8 and n = 48.



Fig. 263: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 8 and n = 48.









with m = 8 and n = 48

min=0.57us max=8.04us avg=1.35us median=1.31us stdev=0.28us



Fig. 266: Histogram of release latency under GSN-EDF-DI, with m=8and n = 48.



Fig. 267: Histogram of tick overhead under GSN-EDF-DI, with m = 8and n = 48.



Fig. 268: Histogram of context-switching overhead under GSN-EDF-DI, with m = 8 and n = 48.



Fig. 269: Histogram of IPI latency under GSN-EDF-DI, with m = 8and n = 48.

min=0.24us max=7.07us avg=1.07us median=1.02us stdev=0.42us samples: total=331039 fraction of samples 0. 0.01 0.00 0.0001 0.20 1.00 1.80 2.60 3.40 4.20 5.00 5.80 6.60 overhead in microseconds (bin size = 0.40us)

Fig. 270: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 8 and n = 48.









min=0.28us max=4.92us avg=0.76us median=0.69us stdev=0.28us







Fig. 274: Histogram of tick overhead under G-EDF-MP, with m = 8 and n = 48.



Fig. 275: Histogram of context-switching overhead under G-EDF-MP, with m = 8 and n = 48.



Fig. 276: Histogram of IPI latency under G-EDF-MP, with m = 8 and n = 48.



Fig. 277: Histogram of client request latency under G-EDF-MP, with m=8 and n=48.

samples: total=291321 fraction of samples 0.1 0.01 0.00 0.0001 0.25 1.25 2.25 3.25 4.25 5.25 6.25 7.25 8.25 overhead in microseconds (bin size = 0.50us)

Fig. 278: Histogram of DSP handler overhead under G-EDF-MP, with m = 8 and n = 48.









min=0.22us max=13.02us avg=1.66us median=1.42us stdev=0.91us samples: total=278505 fraction of samples 0. 0.01 0.001 0.0001 0.40 2.00 3.60 5.20 6.80 8 40 10.00 11.60 13.20 overhead in microseconds (bin size = 0.80us)

Fig. 281: Histogram of release overhead under SD, with m = 8 and n = 48.

min=0.07us max=8.24us avg=1.24us median=1.11us stdev=0.67us



Fig. 282: Histogram of release latency under SD, with m = 8 and n = 48.



Fig. 283: Histogram of tick overhead under SD, with m = 8 and n = 48.



Fig. 284: Histogram of context-switching overhead under SD, with m = 8 and n = 48.



Fig. 285: Histogram of IPI latency under SD, with m = 8 and n = 48.

min=0.24us max=27.88us avg=3.11us median=2.66us stdev=2.16us samples: total=380600 fraction of samples 0.1 0.01 0.001 0.0001 1.00 5.00 9.00 13.00 17.00 21.00 25.00



min=0.07us max=1.26us avg=0.08us median=0.07us stdev=0.05us samples: total=380663 fraction of samples 0.1 0.01 0.001 0.0001 0.05 0.25 0.45 0.65 0.85 1.05 1.25 overhead in microseconds (bin size = 0.10us)

Fig. 287: Histogram of scheduling overhead (b) under GSN-EDF, with m = 8 and n = 56.





samples: total=84295 fraction of samples 0.1 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 4.35 overhead in microseconds (bin size = 0.30us)



min=0.28us max=4.36us avg=0.76us median=0.72us stdev=0.26us



Fig. 290: Histogram of tick overhead under GSN-EDF, with m = 8 and n = 56.



Fig. 291: Histogram of context-switching overhead under GSN-EDF, with m = 8 and n = 56.



Fig. 292: Histogram of IPI latency under GSN-EDF, with m = 8 and n = 56.



Fig. 293: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 8 and n = 56.

min=0.07us max=1.59us avg=0.08us median=0.07us stdev=0.04us samples: total=401600 fraction of samples 0.1 0.01 0.001 0.0001 0.05 0.25 0.45 0.65 0.85 1.05 1.25 1.45

Fig. 294: Histogram of scheduling overhead (b) under GSN-EDF-DI,



Fig. 295: Histogram of release overhead under GSN-EDF-DI, with m = 8 and n = 56.





min=0.13us max=1.62us avg=0.25us median=0.23us stdev=0.07us samples: total=534885 fraction of samples 0.1 0.01 0.001 0.0001 0.05 0.25 0.45 0.65 0.85 1 05 1.25 1.45 1.65 overhead in microseconds (bin size = 0.10us)



overhead in microseconds (bin size = 0.10us) with m = 8 and n = 56.



Fig. 298: Histogram of context-switching overhead under GSN-EDF-DI, with m = 8 and n = 56.



Fig. 299: Histogram of IPI latency under GSN-EDF-DI, with m = 8and n = 56.



Fig. 300: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 8 and n = 56.



Fig. 301: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 8 and n = 56.

min=0.45us max=16.59us avg=4.39us median=3.76us stdev=2.93us 1 samples: total=84291 fraction of samples 0.1 0.01 0.00 0.0001 0.50 2.50 4.50 6.50 8.50 10.50 12.50 14.50 16.50 overhead in microseconds (bin size = 1.00us)

Fig. 302: Histogram of release overhead under G-EDF-MP, with m = 8and n = 56.









min=0.27us max=5.65us avg=1.27us median=1.35us stdev=0.41us samples: total=402137 fraction of samples 0. 0.01 0.001 0.0001 0.20 1.00 1.80 2.60 3.40 4.20 5.00 5.80 overhead in microseconds (bin size = 0.40us)





Fig. 306: Histogram of IPI latency under G-EDF-MP, with m=8 and n=56.



Fig. 307: Histogram of client request latency under G-EDF-MP, with m=8 and n=56.



Fig. 308: Histogram of DSP handler overhead under G-EDF-MP, with m=8 and n=56.



Fig. 309: Histogram of scheduling overhead (a) under SD, with m=8 and n=56.

Fig. 310: Histogram of scheduling overhead (b) under SD, with m = 8 and n = 56.



Fig. 311: Histogram of release overhead under SD, with m = 8 and n = 56.







Fig. 313: Histogram of tick overhead under SD, with m=8 and n=56.

min=0.06us max=23.44us avg=0.89us median=0.09us stdev=1.40us



Fig. 314: Histogram of context-switching overhead under SD, with m = 8 and n = 56.



Fig. 315: Histogram of IPI latency under SD, with m = 8 and n = 56.



Fig. 316: Histogram of scheduling overhead (a) under GSN-EDF, with m = 8 and n = 64.



Fig. 317: Histogram of scheduling overhead (b) under GSN-EDF, with m=8 and n=64.



Fig. 318: Histogram of release overhead under GSN-EDF, with m=8 and n=64.



Fig. 319: Histogram of release latency under GSN-EDF, with m = 8 and n = 64.



Fig. 320: Histogram of tick overhead under GSN-EDF, with m = 8 and n = 64.



Fig. 321: Histogram of context-switching overhead under GSN-EDF, with m = 8 and n = 64.



Fig. 322: Histogram of IPI latency under GSN-EDF, with m = 8 and n = 64.



Fig. 323: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m=8 and n=64.



Fig. 324: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 8 and n = 64.



Fig. 325: Histogram of release overhead under GSN-EDF-DI, with m=8 and n=64.



Fig. 326: Histogram of release latency under GSN-EDF-DI, with m=8 and n=64.



Fig. 327: Histogram of tick overhead under GSN-EDF-DI, with m=8 and n=64.



Fig. 328: Histogram of context-switching overhead under GSN-EDF-DI, with m = 8 and n = 64.



Fig. 329: Histogram of IPI latency under GSN-EDF-DI, with m=8 and n=64.



Fig. 330: Histogram of scheduling overhead (a) under G-EDF-MP, with m=8 and n=64.



Fig. 331: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 8 and n = 64.



Fig. 332: Histogram of release overhead under G-EDF-MP, with m = 8 and n = 64.



Fig. 333: Histogram of release latency under G-EDF-MP, with m=8 and n=64.

min=0.13us max=4.73us avg=0.24us median=0.22us stdev=0.08us samples: total=464337 fraction of samples 0.1 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 4.35 overhead in microseconds (bin size = 0.30us)

Fig. 334: Histogram of tick overhead under G-EDF-MP, with m = 8 and n = 64.















Fig. 338: Histogram of DSP handler overhead under G-EDF-MP, with m=8 and n=64.



Fig. 339: Histogram of scheduling overhead (a) under SD, with m = 8 and n = 64.



Fig. 340: Histogram of scheduling overhead (b) under SD, with m = 8 and n = 64.



Fig. 341: Histogram of release overhead under SD, with m = 8 and n = 64.

min=0.32us max=27.17us avg=1.54us median=0.74us stdev=1.87us 0.01 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.00 0.001 0.00 0.001 0.00 0.001 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.000 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00

Fig. 342: Histogram of release latency under SD, with m = 8 and n = 64.



Fig. 343: Histogram of tick overhead under SD, with m = 8 and n = 64.











Fig. 346: Histogram of scheduling overhead (a) under GSN-EDF, with m = 8 and n = 72.



Fig. 347: Histogram of scheduling overhead (b) under GSN-EDF, with m = 8 and n = 72.



Fig. 348: Histogram of release overhead under GSN-EDF, with m = 8 and n = 72.



Fig. 349: Histogram of release latency under GSN-EDF, with m = 8 and n = 72.



min=0.36us max=5.68us avg=1.26us median=1.37us stdev=0.40us samples: total=473489 fraction of samples 0.1 0.01 0.001 0.0001 0.20 1.00 1.80 2.60 3.40 4.20 5.00 5.80 overhead in microseconds (bin size = 0.40us)













Fig. 354: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 8 and n = 72.



Fig. 355: Histogram of release overhead under GSN-EDF-DI, with m = 8 and n = 72.



Fig. 356: Histogram of release latency under GSN-EDF-DI, with m=8 and n=72.



Fig. 357: Histogram of tick overhead under GSN-EDF-DI, with m = 8 and n = 72.

min=0.35us max=5.71us avg=1.21us median=1.30us stdev=0.41us 1 samples: total=501517 fraction of samples 0.1 0.01 0.001 0.0001 0.20 1.00 1.80 2.60 3.40 4.20 5.00 5.80 overhead in microseconds (bin size = 0.40us)

Fig. 358: Histogram of context-switching overhead under GSN-EDF-DI, with m = 8 and n = 72.









min=0.07us max=1.24us avg=0.08us median=0.07us stdev=0.05us amples: total=504304 0.01 0.001 0.001 0.05 0.25 0.45 0.65 0.85 1.05 1.25 overhead in microseconds (bin size = 0.10us)





Fig. 362: Histogram of release overhead under G-EDF-MP, with m=8 and n=72.



Fig. 363: Histogram of release latency under G-EDF-MP, with m = 8 and n = 72.



Fig. 364: Histogram of tick overhead under G-EDF-MP, with m = 8 and n = 72.



Fig. 365: Histogram of context-switching overhead under G-EDF-MP, with m=8 and n=72.



Fig. 366: Histogram of IPI latency under G-EDF-MP, with m = 8 and n = 72.



Fig. 367: Histogram of client request latency under G-EDF-MP, with m = 8 and n = 72.











Fig. 370: Histogram of scheduling overhead (b) under SD, with m = 8 and n = 72.



Fig. 371: Histogram of release overhead under SD, with m = 8 and n = 72.



Fig. 372: Histogram of release latency under SD, with m = 8 and n = 72.



Fig. 373: Histogram of tick overhead under SD, with m = 8 and n = 72.



Fig. 374: Histogram of context-switching overhead under SD, with m = 8 and n = 72.



Fig. 375: Histogram of IPI latency under SD, with m = 8 and n = 72.









Fig. 378: Histogram of release overhead under GSN-EDF, with m=8 and n=80.



Fig. 379: Histogram of release latency under GSN-EDF, with m = 8 and n = 80.



Fig. 380: Histogram of tick overhead under GSN-EDF, with m = 8 and n = 80.



Fig. 381: Histogram of context-switching overhead under GSN-EDF, with m = 8 and n = 80.



Fig. 382: Histogram of IPI latency under GSN-EDF, with m = 8 and n = 80.







Fig. 384: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 8 and n = 80.



Fig. 385: Histogram of release overhead under GSN-EDF-DI, with m=8 and n=80.



Fig. 386: Histogram of release latency under GSN-EDF-DI, with m=8 and n=80.



Fig. 387: Histogram of tick overhead under GSN-EDF-DI, with m = 8 and n = 80.



Fig. 388: Histogram of context-switching overhead under GSN-EDF-DI, with m = 8 and n = 80.



Fig. 389: Histogram of IPI latency under GSN-EDF-DI, with m = 8 and n = 80.

0.01 0.001 0.001 0.25 1.25 2.25 3.25 4.25 5.25 6.25 7.25 8.25

Fig. 390: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 8 and n = 80.

overhead in microseconds (bin size = 0.50us)



Fig. 391: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 8 and n = 80.



Fig. 392: Histogram of release overhead under G-EDF-MP, with m = 8 and n = 80.



Fig. 393: Histogram of release latency under G-EDF-MP, with m=8 and n=80.

min=0.24us max=8.58us avg=1.08us median=1.03us stdev=0.42us



Fig. 394: Histogram of tick overhead under G-EDF-MP, with m = 8 and n = 80.



Fig. 395: Histogram of context-switching overhead under G-EDF-MP, with m = 8 and n = 80.



Fig. 396: Histogram of IPI latency under G-EDF-MP, with m = 8 and n = 80.



Fig. 397: Histogram of client request latency under G-EDF-MP, with m=8 and n=80.

min=0.07us max=7.61us avg=1.31us median=1.13us stdev=0.76us samples: total=467041 0.01 0.001 0.001 0.25 1.25 2.25 3.25 4.25 5.25 6.25 7.25

overhead in microseconds (bin size = 0.50us) Fig. 398: Histogram of DSP handler overhead under G-EDF-MP, with m = 8 and n = 80.



Fig. 399: Histogram of scheduling overhead (a) under SD, with m = 8 and n = 80.





min=0.23us max=12.87us avg=1.71us median=1.45us stdev=0.92us samples: total=437436 fraction of samples 0. 0.01 0.001 0.0001 0.40 2.00 3.60 5.20 6.80 8 40 10.00 11.60 13.20 overhead in microseconds (bin size = 0.80us)

Fig. 401: Histogram of release overhead under SD, with m = 8 and n = 80.



Fig. 402: Histogram of release latency under SD, with m = 8 and n = 80.



Fig. 403: Histogram of tick overhead under SD, with m = 8 and n = 80.



Fig. 404: Histogram of context-switching overhead under SD, with m = 8 and n = 80.



Fig. 405: Histogram of IPI latency under SD, with m = 8 and n = 80.

min=0.25us max=54.44us avg=2.60us median=1.79us stdev=2.77us samples: total=99546 fraction of samples 0.1 0.01 0.001 0.0001 2.00 10.00 18.00 26.00 34.00 42.00 50.00

overhead in microseconds (bin size = 4.00us) Fig. 406: Histogram of scheduling overhead (a) under GSN-EDF, with m = 16 and n = 16.



Fig. 407: Histogram of scheduling overhead (b) under GSN-EDF, with m = 16 and n = 16.





min=0.30us max=5.73us avg=1.01us median=0.76us stdev=0.62us samples: total=54871 fraction of samples 0. 0.01 0.001 0.0001 0.20 1 00 1.80 2.60 3.40 4.20 5.00 5.80 overhead in microseconds (bin size = 0.40us)



min=0.07us max=1.25us avg=0.09us median=0.07us stdev=0.06us



Fig. 410: Histogram of tick overhead under GSN-EDF, with m = 16 and n = 16.



Fig. 411: Histogram of context-switching overhead under GSN-EDF, with m = 16 and n = 16.



Fig. 412: Histogram of IPI latency under GSN-EDF, with m = 16 and n = 16.



Fig. 413: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 16 and n = 16.

min=0.07us max=1.45us avg=0.09us median=0.07us stdev=0.07us



Fig. 414: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 16 and n = 16.



Fig. 415: Histogram of release overhead under GSN-EDF-DI, with m = 16 and n = 16.





min=0.13us max=4.90us avg=0.37us median=0.32us stdev=0.22us samples: total=1379314 fraction of samples 0. 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 4.35 4.95 overhead in microseconds (bin size = 0.30us)





Fig. 418: Histogram of context-switching overhead under GSN-EDF-DI, with m = 16 and n = 16.



Fig. 419: Histogram of IPI latency under GSN-EDF-DI, with m = 16 and n = 16.



Fig. 420: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 16 and n = 16.



Fig. 421: Histogram of scheduling overhead (b) under G-EDF-MP, with m=16 and n=16.

min=0.54us max=18.52us avg=2.80us median=2.17us stdev=2.03us



Fig. 422: Histogram of release overhead under G-EDF-MP, with m = 16 and n = 16.



Fig. 423: Histogram of release latency under G-EDF-MP, with m = 16 and n = 16.





min=0.35us max=5.44us avg=1.57us median=1.58us stdev=0.54us sdmples: total=99582 0.01 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.00 1.80 2.60 3.40 4.20 5.00 overhead in microseconds (bin size = 0.40us)





Fig. 426: Histogram of IPI latency under G-EDF-MP, with m=16 and n=16.



Fig. 427: Histogram of client request latency under G-EDF-MP, with m = 16 and n = 16.



Fig. 428: Histogram of DSP handler overhead under G-EDF-MP, with m = 16 and n = 16.



Fig. 429: Histogram of scheduling overhead (a) under SD, with m=16 and n=16.



Fig. 430: Histogram of scheduling overhead (b) under SD, with m = 16 and n = 16.



Fig. 431: Histogram of release overhead under SD, with m = 16 and n = 16.







Fig. 433: Histogram of tick overhead under SD, with m = 16 and n = 16.



Fig. 434: Histogram of context-switching overhead under SD, with m = 16 and n = 16.



Fig. 435: Histogram of IPI latency under SD, with m = 16 and n = 16.



Fig. 436: Histogram of scheduling overhead (a) under GSN-EDF, with m = 16 and n = 32.



Fig. 437: Histogram of scheduling overhead (b) under GSN-EDF, with m = 16 and n = 32.

samples: total=67744 fraction of samples 0.1 0.01 0.001 0.0001 2.00 10.00 18.00 26.00 34.00 42.00 50.00 58.00 66.00 overhead in microseconds (bin size = 4.00us)

Fig. 438: Histogram of release overhead under GSN-EDF, with m = 16and n = 32.



overhead in microseconds (bin size = 0.50us) Fig. 439: Histogram of release latency under GSN-EDF, with m = 16and n = 32.









min=0.36us max=64.10us avg=6.68us median=4.56us stdev=6.30us



Fig. 442: Histogram of IPI latency under GSN-EDF, with m = 16 and n = 32.



Fig. 443: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 16 and n = 32.



Fig. 444: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 16 and n = 32.



Fig. 445: Histogram of release overhead under GSN-EDF-DI, with m = 16 and n = 32.



Fig. 446: Histogram of release latency under GSN-EDF-DI, with m =16 and n = 32.









min=0.80us max=11.76us avg=1.96us median=1.96us stdev=0.57us samples: total=165465 fraction of samples 0.1 0.01 0.001 0.0001 0.35 1.75 3.15 4.55 5.95 7.35 8.75 10.15 11.55 overhead in microseconds (bin size = 0.70us)





Fig. 450: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 16 and n = 32.



Fig. 451: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 16 and n = 32.



Fig. 452: Histogram of release overhead under G-EDF-MP, with m =16 and n = 32.



Fig. 453: Histogram of release latency under G-EDF-MP, with m = 16and n = 32.

min=0.13us max=5.03us avg=0.35us median=0.29us stdev=0.20us samples: total=1121452 fraction of samples 0.1 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 4.35 4.95 overhead in microseconds (bin size = 0.30us)

Fig. 454: Histogram of tick overhead under G-EDF-MP, with m = 16and n = 32.













overhead in microseconds (bin size = 0.60us)



Fig. 458: Histogram of DSP handler overhead under G-EDF-MP, with m = 16 and n = 32.



Fig. 459: Histogram of scheduling overhead (a) under SD, with m = 16 and n = 32.



Fig. 460: Histogram of scheduling overhead (b) under SD, with m = 16 and n = 32.



Fig. 461: Histogram of release overhead under SD, with m = 16 and n = 32.



Fig. 462: Histogram of release latency under SD, with m = 16 and n = 32.



Fig. 463: Histogram of tick overhead under SD, with m = 16 and n = 32.











Fig. 466: Histogram of scheduling overhead (a) under GSN-EDF, with m = 16 and n = 48.



Fig. 467: Histogram of scheduling overhead (b) under GSN-EDF, with m = 16 and n = 48.



Fig. 468: Histogram of release overhead under GSN-EDF, with m = 16and n = 48.



Fig. 469: Histogram of release latency under GSN-EDF, with m = 16and n = 48.

samples: total=1077198 fraction of samples 0.1 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 4.35 overhead in microseconds (bin size = 0.30us)

Fig. 470: Histogram of tick overhead under GSN-EDF, with m = 16and n = 48.



overhead in microseconds (bin size = 0.40us) Fig. 471: Histogram of context-switching overhead under GSN-EDF, with m = 16 and n = 48.





min=0.35us max=80.72us avg=9.33us median=6.65us stdev=8.16us samples: total=315974 fraction of samples 0.1 0.01 0.001 0.0001 2.50 12.50 22.50 32.50 42.50 52.50 62.50 72.50 82.50 overhead in microseconds (bin size = 5.00us)



min=0.40us max=6.02us avg=1.83us median=1.86us stdev=0.45us




Fig. 474: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 16 and n = 48.



Fig. 475: Histogram of release overhead under GSN-EDF-DI, with m = 16 and n = 48.



Fig. 476: Histogram of release latency under GSN-EDF-DI, with m = 16 and n = 48.



Fig. 477: Histogram of tick overhead under GSN-EDF-DI, with m=16 and n=48.



Fig. 478: Histogram of context-switching overhead under GSN-EDF-DI, with m = 16 and n = 48.









min=0.07us max=1.60us avg=0.09us median=0.07us stdev=0.07us samples: total=316136 fraction of samples 0.1 0.01 0.001 0.0001 0.05 0.25 0.45 0.65 0.85 1.05 1.25 1.45 1.65 overhead in microseconds (bin size = 0.10us)





Fig. 482: Histogram of release overhead under G-EDF-MP, with m = 16 and n = 48.



Fig. 483: Histogram of release latency under G-EDF-MP, with m = 16 and n = 48.



Fig. 484: Histogram of tick overhead under G-EDF-MP, with m = 16 and n = 48.



Fig. 485: Histogram of context-switching overhead under G-EDF-MP, with m = 16 and n = 48.



Fig. 486: Histogram of IPI latency under G-EDF-MP, with m = 16 and n = 48.



Fig. 487: Histogram of client request latency under G-EDF-MP, with m = 16 and n = 48.



Fig. 488: Histogram of DSP handler overhead under G-EDF-MP, with m = 16 and n = 48.



Fig. 489: Histogram of scheduling overhead (a) under SD, with m = 16 and n = 48.



Fig. 490: Histogram of scheduling overhead (b) under SD, with m = 16 and n = 48.



Fig. 491: Histogram of release overhead under SD, with m = 16 and n = 48.



Fig. 492: Histogram of release latency under SD, with m = 16 and n = 48.



Fig. 493: Histogram of tick overhead under SD, with m = 16 and n = 48.



Fig. 494: Histogram of context-switching overhead under SD, with m = 16 and n = 48.



Fig. 495: Histogram of IPI latency under SD, with m = 16 and n = 48.





min=0.07us max=1.75us avg=0.08us median=0.07us stdev=0.05us samples: total=412095 raction of samples 0.1 0.01 0.00 0 0001 0.05 0.25 0.45 0.65 0.85 1.25 1.45 1.05 1.65 overhead in microseconds (bin size = 0.10us)





Fig. 498: Histogram of release overhead under GSN-EDF, with m=16 and n=64.



Fig. 499: Histogram of release latency under GSN-EDF, with m = 16 and n = 64.



Fig. 500: Histogram of tick overhead under GSN-EDF, with m = 16 and n = 64.



Fig. 501: Histogram of context-switching overhead under GSN-EDF, with m = 16 and n = 64.





min=0.33us max=75.07us avg=11.41us median=8.37us stdev=9.35us



Fig. 503: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 16 and n = 64.







Fig. 505: Histogram of release overhead under GSN-EDF-DI, with m = 16 and n = 64.



Fig. 506: Histogram of release latency under GSN-EDF-DI, with m =16 and n = 64.



Fig. 507: Histogram of tick overhead under GSN-EDF-DI, with m = 16and n = 64.



Fig. 508: Histogram of context-switching overhead under GSN-EDF-DI, with m = 16 and n = 64.



Fig. 509: Histogram of IPI latency under GSN-EDF-DI, with m = 16and n = 64.



Fig. 510: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 16 and n = 64.



Fig. 511: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 16 and n = 64.





min=0.29us max=6.87us avg=0.92us median=0.72us stdev=0.59us samples: total=85976 fraction of samples 0.1 0.01 0.00 0.0001 0.20 1.00 1.80 2.60 3.40 5.00 5.80 6.60 4.20 overhead in microseconds (bin size = 0.40us)





Fig. 514: Histogram of tick overhead under G-EDF-MP, with m = 16 and n = 64.



Fig. 515: Histogram of context-switching overhead under G-EDF-MP, with m = 16 and n = 64.



Fig. 516: Histogram of IPI latency under G-EDF-MP, with m = 16 and n = 64.



Fig. 517: Histogram of client request latency under G-EDF-MP, with m=16 and n=64.

min=0.07us max=10.22us avg=1.47us median=1.34us stdev=0.75us samples: total=398767 fraction of samples 0. 0.01 0.00 0.0001 0.30 1.50 2.70 3.90 5.10 6.30 7.50 8.70 9.90 overhead in microseconds (bin size = 0.60us)

Fig. 518: Histogram of DSP handler overhead under G-EDF-MP, with m = 16 and n = 64.











Fig. 521: Histogram of release overhead under SD, with m = 16 and n = 64.



Fig. 522: Histogram of release latency under SD, with m = 16 and n = 64.



Fig. 523: Histogram of tick overhead under SD, with m = 16 and n = 64.



Fig. 524: Histogram of context-switching overhead under SD, with m = 16 and n = 64.



Fig. 525: Histogram of IPI latency under SD, with m = 16 and n = 64.

min=0.27us max=68.47us avg=10.99us median=7.02us stdev=10.17us samples: total=496913 fraction of samples 0.1 0.01 0.00 0.0001 2.00 10.00 18.00 26.00 34.00 42.00 50.00 58.00 66.00 overhead in microseconds (bin size = 4.00us)



min=0.07us max=1.55us avg=0.08us median=0.07us stdev=0.05us samples: total=496998 0.1 0.01 0.001 0.0001 0.05 0.25 0.45 0.65 0.85 1.05 1.25 1.45 overhead in microseconds (bin size = 0.10us)

Fig. 527: Histogram of scheduling overhead (b) under GSN-EDF, with m = 16 and n = 80.





min=0.30us max=7.49us avg=0.98us median=0.74us stdev=0.65us samples: total=87676 fraction of samples 0.1 0.01 0.001 0.0001 0.25 1 25 2.25 3 25 4 25 5.25 6.25 7.25 overhead in microseconds (bin size = 0.50us)



fraction of samples



Fig. 530: Histogram of tick overhead under GSN-EDF, with m = 16and n = 80.



Fig. 531: Histogram of context-switching overhead under GSN-EDF, with m = 16 and n = 80.



Fig. 532: Histogram of IPI latency under GSN-EDF, with m = 16 and n = 80.



Fig. 533: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 16 and n = 80.

samples: total=515369 fraction of samples 0.1 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 overhead in microseconds (bin size = 0.30us)

Fig. 534: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 16 and n = 80.



Fig. 535: Histogram of release overhead under GSN-EDF-DI, with m = 16 and n = 80.









min=0.07us max=3.80us avg=0.08us median=0.07us stdev=0.05us



Fig. 538: Histogram of context-switching overhead under GSN-EDF-DI, with m = 16 and n = 80.



Fig. 539: Histogram of IPI latency under GSN-EDF-DI, with m = 16 and n = 80.



Fig. 540: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 16 and n = 80.



Fig. 541: Histogram of scheduling overhead (b) under G-EDF-MP, with m=16 and n=80.

min=0.43us max=40.53us avg=7.51us median=5.29us stdev=6.61us 0.01 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.00

Fig. 542: Histogram of release overhead under G-EDF-MP, with m = 16 and n = 80.



Fig. 543: Histogram of release latency under G-EDF-MP, with m = 16 and n = 80.







Fig. 545: Histogram of context-switching overhead under G-EDF-MP, with m=16 and n=80.



Fig. 546: Histogram of IPI latency under G-EDF-MP, with m = 16 and n = 80.



Fig. 547: Histogram of client request latency under G-EDF-MP, with m = 16 and n = 80.



Fig. 548: Histogram of DSP handler overhead under G-EDF-MP, with m = 16 and n = 80.



Fig. 549: Histogram of scheduling overhead (a) under SD, with m=16 and n=80.

Fig. 550: Histogram of scheduling overhead (b) under SD, with m = 16 and n = 80.



Fig. 551: Histogram of release overhead under SD, with m = 16 and n = 80.







Fig. 553: Histogram of tick overhead under SD, with m = 16 and n = 80.

min=0.06us max=74.13us avg=2.29us median=0.09us stdev=4.21us



Fig. 554: Histogram of context-switching overhead under SD, with m = 16 and n = 80.



Fig. 555: Histogram of IPI latency under SD, with m = 16 and n = 80.



Fig. 556: Histogram of scheduling overhead (a) under GSN-EDF, with m = 16 and n = 96.



Fig. 557: Histogram of scheduling overhead (b) under GSN-EDF, with m=16 and n=96.

min=0.45us max=79.73us avg=15.16us median=10.94us stdev=12.63us samples: total=91577 fraction of samples 0.1 0.01 0.001 0.0001 2.50 12.50 22.50 32.50 42.50 52.50 62.50 72.50 overhead in microseconds (bin size = 5.00us)

Fig. 558: Histogram of release overhead under GSN-EDF, with m = 16 and n = 96.





Fig. 559: Histogram of release latency under GSN-EDF, with m = 16 and n = 96.



Fig. 560: Histogram of tick overhead under GSN-EDF, with m = 16 and n = 96.



Fig. 561: Histogram of context-switching overhead under GSN-EDF, with m = 16 and n = 96.



Fig. 562: Histogram of IPI latency under GSN-EDF, with m = 16 and n = 96.



Fig. 563: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 16 and n = 96.



Fig. 564: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 16 and n = 96.



Fig. 565: Histogram of release overhead under GSN-EDF-DI, with m = 16 and n = 96.

1 samples: total=91583 fraction of samples 0.1 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 4.35 overhead in microseconds (bin size = 0.30us)

Fig. 566: Histogram of release latency under GSN-EDF-DI, with m =16 and n = 96.



Fig. 567: Histogram of tick overhead under GSN-EDF-DI, with m = 16and n = 96.









min=0.29us max=4.52us avg=1.04us median=0.79us stdev=0.64us



Fig. 570: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 16 and n = 96.



Fig. 571: Histogram of scheduling overhead (b) under G-EDF-MP, with m=16 and n=96.



Fig. 572: Histogram of release overhead under G-EDF-MP, with m = 16 and n = 96.



Fig. 573: Histogram of release latency under G-EDF-MP, with m=16 and n=96.

min=0.13us max=4.43us avg=0.34us median=0.26us stdev=0.20us samples: total=980089 fraction of samples 0.1 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 4.35 overhead in microseconds (bin size = 0.30us)

Fig. 574: Histogram of tick overhead under G-EDF-MP, with m = 16 and n = 96.



Fig. 575: Histogram of context-switching overhead under G-EDF-MP, with m = 16 and n = 96.











Fig. 578: Histogram of DSP handler overhead under G-EDF-MP, with m = 16 and n = 96.



Fig. 579: Histogram of scheduling overhead (a) under SD, with m = 16and n = 96.



Fig. 580: Histogram of scheduling overhead (b) under SD, with m = 16and n = 96.



Fig. 581: Histogram of release overhead under SD, with m = 16 and n = 96.

min=0.30us max=102.48us avg=3.50us median=0.81us stdev=7.30us samples: total=523709 fraction of samples 0.1 0.01 0.001 0.0001 105.00 5.00 65.00 85.00 25.00 45.00 overhead in microseconds (bin size = 10.00us)





Fig. 583: Histogram of tick overhead under SD, with m = 16 and n = 96.











Fig. 586: Histogram of scheduling overhead (a) under GSN-EDF, with m = 16 and n = 112.



Fig. 587: Histogram of scheduling overhead (b) under GSN-EDF, with m=16 and n=112.



Fig. 588: Histogram of release overhead under GSN-EDF, with m = 16 and n = 112.



Fig. 589: Histogram of release latency under GSN-EDF, with m=16 and n=112.

min=0.13us max=4.03us avg=0.42us median=0.39us stdev=0.22us amples: total=1084536 0.01 0.001 0.001 0.001 0.001 0.001 0.05 0.75 1.35 1.95 2.55 3.15 3.75 overhead in microseconds (bin size = 0.30us)





Fig. 591: Histogram of context-switching overhead under GSN-EDF, with m = 16 and n = 112.









min=0.36us max=6.52us avg=1.70us median=1.74us stdev=0.45us



Fig. 594: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 16 and n = 112.



Fig. 595: Histogram of release overhead under GSN-EDF-DI, with m = 16 and n = 112.



Fig. 596: Histogram of release latency under GSN-EDF-DI, with m = 16 and n = 112.



Fig. 597: Histogram of tick overhead under GSN-EDF-DI, with m=16 and n=112.



Fig. 598: Histogram of context-switching overhead under GSN-EDF-DI, with m = 16 and n = 112.







Fig. 600: Histogram of scheduling overhead (a) under G-EDF-MP, with m=16 and n=112.







Fig. 602: Histogram of release overhead under G-EDF-MP, with m = 16 and n = 112.



Fig. 603: Histogram of release latency under G-EDF-MP, with m = 16 and n = 112.



Fig. 604: Histogram of tick overhead under G-EDF-MP, with m = 16 and n = 112.



Fig. 605: Histogram of context-switching overhead under G-EDF-MP, with m=16 and n=112.



Fig. 606: Histogram of IPI latency under G-EDF-MP, with m = 16 and n = 112.



Fig. 607: Histogram of client request latency under G-EDF-MP, with m=16 and n=112.











Fig. 610: Histogram of scheduling overhead (b) under SD, with m = 16and n = 112.



Fig. 611: Histogram of release overhead under SD, with m = 16 and n = 112.



Fig. 612: Histogram of release latency under SD, with m = 16 and n = 112.



Fig. 613: Histogram of tick overhead under SD, with m = 16 and n = 112.

min=0.26us max=6.08us avg=1.59us median=1.63us stdev=0.51us 1 samples: total=779228 fraction of samples 0.1 0.01 0.001 0.0001 0.20 1.00 1.80 2.60 3.40 4.20 5.00 5.80 overhead in microseconds (bin size = 0.40us)









Fig. 616: Histogram of scheduling overhead (a) under GSN-EDF, with m = 16 and n = 128.





112.



Fig. 618: Histogram of release overhead under GSN-EDF, with m = 16and n = 128.



Fig. 619: Histogram of release latency under GSN-EDF, with m = 16and n = 128.



Fig. 620: Histogram of tick overhead under GSN-EDF, with m = 16and n = 128.



Fig. 621: Histogram of context-switching overhead under GSN-EDF, with m = 16 and n = 128.



Fig. 622: Histogram of IPI latency under GSN-EDF, with m = 16 and n = 128.



Fig. 623: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 16 and n = 128.







Fig. 625: Histogram of release overhead under GSN-EDF-DI, with m = 16 and n = 128.



Fig. 626: Histogram of release latency under GSN-EDF-DI, with m = 16 and n = 128.



Fig. 627: Histogram of tick overhead under GSN-EDF-DI, with m = 16 and n = 128.



Fig. 628: Histogram of context-switching overhead under GSN-EDF-DI, with m = 16 and n = 128.



Fig. 629: Histogram of IPI latency under GSN-EDF-DI, with m = 16 and n = 128.

samples: total=877818 fraction of samples 0. 0.01 0.00 0.0001 0.35 1.75 3.15 4.55 5.95 7.35 8.75 10.15 11.55 overhead in microseconds (bin size = 0.70us)

Fig. 630: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 16 and n = 128.



Fig. 631: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 16 and n = 128.



Fig. 632: Histogram of release overhead under G-EDF-MP, with m = 16 and n = 128.



Fig. 633: Histogram of release latency under G-EDF-MP, with m = 16 and n = 128.

min=0.19us max=11.53us avg=1.46us median=1.42us stdev=0.50us



Fig. 634: Histogram of tick overhead under G-EDF-MP, with m = 16 and n = 128.



Fig. 635: Histogram of context-switching overhead under G-EDF-MP, with m = 16 and n = 128.



Fig. 636: Histogram of IPI latency under G-EDF-MP, with m = 16 and n = 128.



Fig. 637: Histogram of client request latency under G-EDF-MP, with m = 16 and n = 128.

min=0.07us max=10.50us avg=1.69us median=1.44us stdev=0.95us samples: total=761756 fraction of samples 0. 0.0 0.00 0.0001 0.30 1.50 2.70 3.90 5.10 6.30 7.50 8.70 9.90

overhead in microseconds (bin size = 0.60us) Fig. 638: Histogram of DSP handler overhead under G-EDF-MP, with m = 16 and n = 128.











Fig. 641: Histogram of release overhead under SD, with m = 16 and n = 128.



Fig. 642: Histogram of release latency under SD, with m = 16 and n = 128.



Fig. 643: Histogram of tick overhead under SD, with m = 16 and n = 128.



Fig. 644: Histogram of context-switching overhead under SD, with m = 16 and n = 128.



Fig. 645: Histogram of IPI latency under SD, with m = 16 and n = 128.

min=0.33us max=86.72us avg=16.02us median=11.55us stdev=12.59us



Fig. 646: Histogram of scheduling overhead (a) under GSN-EDF, with m = 16 and n = 144.



Fig. 647: Histogram of scheduling overhead (b) under GSN-EDF, with m=16 and n=144.











Fig. 650: Histogram of tick overhead under GSN-EDF, with m = 16 and n = 144.



Fig. 651: Histogram of context-switching overhead under GSN-EDF, with m = 16 and n = 144.



Fig. 652: Histogram of IPI latency under GSN-EDF, with m = 16 and n = 144.



Fig. 653: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 16 and n = 144.

min=0.07us max=1.51us avg=0.08us median=0.07us stdev=0.06us



Fig. 654: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 16 and n = 144.



Fig. 655: Histogram of release overhead under GSN-EDF-DI, with m=16 and n=144.





min=0.14us max=4.12us avg=0.44us median=0.41us stdev=0.21us and the state of the s





Fig. 658: Histogram of context-switching overhead under GSN-EDF-DI, with m = 16 and n = 144.



Fig. 659: Histogram of IPI latency under GSN-EDF-DI, with m = 16and n = 144.



Fig. 660: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 16 and n = 144.



Fig. 661: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 16 and n = 144.

min=0.60us max=42.71us avg=13.08us median=11.19us stdev=9.02us samples: total=94070 fraction of samples 0.1 0.01 0.001 0.0001 1.50 7.50 13.50 19.50 25.50 31.50 37.50 43.50 overhead in microseconds (bin size = 3.00us)



min=0.28us max=7.76us avg=0.89us median=0.72us stdev=0.56us samples: total=94075 fraction of samples 0. 0.01 0.001 0.0001 0.25 1.25 2.25 3.25 4.25 5.25 6.25 7.25 overhead in microseconds (bin size = 0.50us)





Fig. 664: Histogram of tick overhead under G-EDF-MP, with m = 16and n = 144.



Fig. 665: Histogram of context-switching overhead under G-EDF-MP, with m = 16 and n = 144.



Fig. 666: Histogram of IPI latency under G-EDF-MP, with m=16 and n=144.



Fig. 667: Histogram of client request latency under G-EDF-MP, with m = 16 and n = 144.



Fig. 668: Histogram of DSP handler overhead under G-EDF-MP, with m = 16 and n = 144.



Fig. 669: Histogram of scheduling overhead (a) under SD, with m=16 and n=144.

Fig. 670: Histogram of scheduling overhead (b) under SD, with m = 16 and n = 144.



Fig. 671: Histogram of release overhead under SD, with m = 16 and n = 144.







Fig. 673: Histogram of tick overhead under SD, with m = 16 and n = 144.

min=0.06us max=110.67us avg=2.43us median=0.08us stdev=5.02us



Fig. 674: Histogram of context-switching overhead under SD, with m = 16 and n = 144.



Fig. 675: Histogram of IPI latency under SD, with m = 16 and n = 144.



Fig. 676: Histogram of scheduling overhead (a) under GSN-EDF, with m = 16 and n = 160.



Fig. 677: Histogram of scheduling overhead (b) under GSN-EDF, with m=16 and n=160.

min=0.67us max=92.01us avg=23.26us median=19.54us stdev=16.18us



Fig. 678: Histogram of release overhead under GSN-EDF, with m = 16 and n = 160.







Fig. 680: Histogram of tick overhead under GSN-EDF, with m = 16 and n = 160.



Fig. 681: Histogram of context-switching overhead under GSN-EDF, with m = 16 and n = 160.



Fig. 682: Histogram of IPI latency under GSN-EDF, with m = 16 and n = 160.



Fig. 683: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 16 and n = 160.



Fig. 684: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 16 and n = 160.



Fig. 685: Histogram of release overhead under GSN-EDF-DI, with m = 16 and n = 160.

min=0.31us max=5.62us avg=1.04us median=0.79us stdev=0.65us samples: total=94369 fraction of samples 0.1 0.01 0.001 0.0001 0.20 1.00 1.80 2.60 3.40 4.20 5.00 5.80 overhead in microseconds (bin size = 0.40us)



samples: total=1009975 fraction of samples 0. 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 4.35 overhead in microseconds (bin size = 0.30us)





Fig. 688: Histogram of context-switching overhead under GSN-EDF-DI, with m = 16 and n = 160.



Fig. 689: Histogram of IPI latency under GSN-EDF-DI, with m = 16and n = 160.

min=0.14us max=4.59us avg=0.44us median=0.41us stdev=0.21us



Fig. 690: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 16 and n = 160.



Fig. 691: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 16 and n = 160.



Fig. 692: Histogram of release overhead under G-EDF-MP, with m =16 and n = 160.



Fig. 693: Histogram of release latency under G-EDF-MP, with m = 16and n = 160.

samples: total=992876 fraction of samples 0.1 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 4.35 overhead in microseconds (bin size = 0.30us)

Fig. 694: Histogram of tick overhead under G-EDF-MP, with m = 16and n = 160.



overhead in microseconds (bin size = 0.40us) Fig. 695: Histogram of context-switching overhead under G-EDF-MP, with m = 16 and n = 160.









min=0.13us max=4.55us avg=0.33us median=0.25us stdev=0.19us



Fig. 698: Histogram of DSP handler overhead under G-EDF-MP, with m = 16 and n = 160.



Fig. 699: Histogram of scheduling overhead (a) under SD, with m = 16 and n = 160.



Fig. 700: Histogram of scheduling overhead (b) under SD, with m = 16 and n = 160.



Fig. 701: Histogram of release overhead under SD, with m = 16 and n = 160.

min=0.28us max=147.08us avg=6.47us median=0.83us stdev=12.26us samples: total=803558 fraction of samples 0.1 0.01 0.001 0.0001 5.00 25.00 45.00 65.00 85.00 105.00 125.00 145.00 overhead in microseconds (bin size = 10.00us)

Fig. 702: Histogram of release latency under SD, with m = 16 and n = 160.











Fig. 705: Histogram of IPI latency under SD, with m = 16 and n = 160.



Fig. 706: Histogram of scheduling overhead (a) under GSN-EDF, with m = 24 and n = 24.



Fig. 707: Histogram of scheduling overhead (b) under GSN-EDF, with m = 24 and n = 24.



Fig. 708: Histogram of release overhead under GSN-EDF, with m=24and n = 24.



Fig. 709: Histogram of release latency under GSN-EDF, with m = 24and n = 24.

min=0.14us max=4.37us avg=0.38us median=0.29us stdev=0.26us samples: total=1675909 fraction of samples 0.1 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 4.35 overhead in microseconds (bin size = 0.30us)





Fig. 711: Histogram of context-switching overhead under GSN-EDF, with m = 24 and n = 24.











Fig. 714: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 24 and n = 24.



Fig. 715: Histogram of release overhead under GSN-EDF-DI, with m = 24 and n = 24.



Fig. 716: Histogram of release latency under GSN-EDF-DI, with m =24 and n = 24.



Fig. 717: Histogram of tick overhead under GSN-EDF-DI, with m = 24and n = 24.



Fig. 718: Histogram of context-switching overhead under GSN-EDF-DI, with m = 24 and n = 24.



Fig. 719: Histogram of IPI latency under GSN-EDF-DI, with m=24and n = 24.





min=0.07us max=2.55us avg=0.11us median=0.07us stdev=0.08us samples: total=139303 fraction of samples 0.1 0.01 0.001 0.0001 0.10 0.50 0.90 1.30 1.70 2.10 2.50 overhead in microseconds (bin size = 0.20us)





Fig. 722: Histogram of release overhead under G-EDF-MP, with m = 24 and n = 24.



Fig. 723: Histogram of release latency under G-EDF-MP, with m = 24 and n = 24.



Fig. 724: Histogram of tick overhead under G-EDF-MP, with m = 24 and n = 24.



Fig. 725: Histogram of context-switching overhead under G-EDF-MP, with m = 24 and n = 24.



Fig. 726: Histogram of IPI latency under G-EDF-MP, with m=24 and n=24.



Fig. 727: Histogram of client request latency under G-EDF-MP, with m = 24 and n = 24.







Fig. 729: Histogram of scheduling overhead (a) under SD, with m = 24 and n = 24.



Fig. 730: Histogram of scheduling overhead (b) under SD, with m = 24 and n = 24.



Fig. 731: Histogram of release overhead under SD, with m = 24 and n = 24.



Fig. 732: Histogram of release latency under SD, with m = 24 and n = 24.



Fig. 733: Histogram of tick overhead under SD, with m = 24 and n = 24.



Fig. 734: Histogram of context-switching overhead under SD, with m = 24 and n = 24.



Fig. 735: Histogram of IPI latency under SD, with m = 24 and n = 24.



Fig. 736: Histogram of scheduling overhead (a) under GSN-EDF, with m=24 and n=48.





m = m



Fig. 738: Histogram of release overhead under GSN-EDF, with m = 24 and n = 48.



Fig. 739: Histogram of release latency under GSN-EDF, with m = 24 and n = 48.



Fig. 740: Histogram of tick overhead under GSN-EDF, with m = 24 and n = 48.



Fig. 741: Histogram of context-switching overhead under GSN-EDF, with m = 24 and n = 48.



Fig. 742: Histogram of IPI latency under GSN-EDF, with m = 24 and n = 48.

min=0.39us max=105.20us avg=11.94us median=8.39us stdev=11.21us



Fig. 743: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 24 and n = 48.



Fig. 744: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 24 and n = 48.



Fig. 745: Histogram of release overhead under GSN-EDF-DI, with m=24 and n=48.



Fig. 746: Histogram of release latency under GSN-EDF-DI, with m = 24 and n = 48.



Fig. 747: Histogram of tick overhead under GSN-EDF-DI, with m = 24 and n = 48.



Fig. 748: Histogram of context-switching overhead under GSN-EDF-DI, with m = 24 and n = 48.



Fig. 749: Histogram of IPI latency under GSN-EDF-DI, with m = 24 and n = 48.

min=0.24us max=11.52us avg=1.61us median=1.54us stdev=0.55us samples: total=303515 fraction of samples 0.1 0.01 0.00 0.0001 0.35 1.75 3.15 4.55 5.95 7.35 8.75 10.15 11.55 overhead in microseconds (bin size = 0.70us)

Fig. 750: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 24 and n = 48.



Fig. 751: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 24 and n = 48.



Fig. 752: Histogram of release overhead under G-EDF-MP, with m = 24 and n = 48.

min=0.33us max=7.14us avg=1.11us median=0.71us stdev=0.94us samples: total=80132 fraction of samples 0.1 0.01 0.001 0.0001 0.20 1.00 1.80 2.60 3.40 4.20 5.00 5.80 6.60 overhead in microseconds (bin size = 0.40us)





Fig. 754: Histogram of tick overhead under G-EDF-MP, with m = 24and n = 48.



Fig. 755: Histogram of context-switching overhead under G-EDF-MP, with m = 24 and n = 48.



Fig. 756: Histogram of IPI latency under G-EDF-MP, with m=24 and n = 48.



Fig. 757: Histogram of client request latency under G-EDF-MP, with m = 24 and n = 48.

min=0.07us max=8.86us avg=1.38us median=1.30us stdev=0.55us samples: total=301099 fraction of samples 0. 0.01 0.00 0.0001 0.25 1.25 2.25 3.25 4.25 5.25 6.25 7.25 8.25
















Fig. 762: Histogram of release latency under SD, with m = 24 and n = 48.



Fig. 763: Histogram of tick overhead under SD, with m = 24 and n = 48.



Fig. 764: Histogram of context-switching overhead under SD, with m = 24 and n = 48.



Fig. 765: Histogram of IPI latency under SD, with m = 24 and n = 48.

min=0.25us max=133.27us avg=16.59us median=10.68us stdev=16.68us



Fig. 766: Histogram of scheduling overhead (a) under GSN-EDF, with m = 24 and n = 72.



Fig. 767: Histogram of scheduling overhead (b) under GSN-EDF, with m = 24 and n = 72.





min=0.29us max=8.33us avg=1.18us median=0.81us stdev=0.97us samples: total=85333 fraction of samples 0. 0.01 0.001 0.0001 0.25 1.25 2 25 3 25 4.25 5.25 6.25 7.25 8.25 overhead in microseconds (bin size = 0.50us)





Fig. 770: Histogram of tick overhead under GSN-EDF, with m = 24 and n = 72.



Fig. 771: Histogram of context-switching overhead under GSN-EDF, with m = 24 and n = 72.



Fig. 772: Histogram of IPI latency under GSN-EDF, with m = 24 and n = 72.



Fig. 773: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 24 and n = 72.

min=0.07us max=1.97us avg=0.08us median=0.07us stdev=0.06us



Fig. 774: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 24 and n = 72.





Fig. 775: Histogram of release overhead under GSN-EDF-DI, with m = 24 and n = 72.











Fig. 778: Histogram of context-switching overhead under GSN-EDF-DI, with m = 24 and n = 72.



Fig. 779: Histogram of IPI latency under GSN-EDF-DI, with m = 24 and n = 72.



Fig. 780: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 24 and n = 72.



Fig. 781: Histogram of scheduling overhead (b) under G-EDF-MP, with m=24 and n=72.

min=0.48us max=52.86us avg=7.49us median=5.00us stdev=7.23us samples: total=85333 fraction of samples 0.1 0.01 0.00 0.0001 1.50 7.50 13.50 19.50 25.50 37.50 43.50 49.50 overhead in microseconds (bin size = 3.00us)

Fig. 782: Histogram of release overhead under G-EDF-MP, with m = 24 and n = 72.





Fig. 783: Histogram of release latency under G-EDF-MP, with m=24 and n=72.



Fig. 784: Histogram of tick overhead under G-EDF-MP, with m = 24 and n = 72.



Fig. 785: Histogram of context-switching overhead under G-EDF-MP, with m=24 and n=72.



Fig. 786: Histogram of IPI latency under G-EDF-MP, with m = 24 and n = 72.



Fig. 787: Histogram of client request latency under G-EDF-MP, with m = 24 and n = 72.



Fig. 788: Histogram of DSP handler overhead under G-EDF-MP, with m = 24 and n = 72.



Fig. 789: Histogram of scheduling overhead (a) under SD, with m = 24 and n = 72.



Fig. 790: Histogram of scheduling overhead (b) under SD, with m = 24 and n = 72.



Fig. 791: Histogram of release overhead under SD, with m = 24 and n = 72.







Fig. 793: Histogram of tick overhead under SD, with m = 24 and n = 72.



Fig. 794: Histogram of context-switching overhead under SD, with m = 24 and n = 72.



Fig. 795: Histogram of IPI latency under SD, with m = 24 and n = 72.



Fig. 796: Histogram of scheduling overhead (a) under GSN-EDF, with m = 24 and n = 96.



Fig. 797: Histogram of scheduling overhead (b) under GSN-EDF, with m=24 and n=96.

min=0.51us max=127.55us avg=20.36us median=13.14us stdev=19.89us



Fig. 798: Histogram of release overhead under GSN-EDF, with m=24 and n=96.

min=0.30us max=7.58us avg=1.20us median=0.83us stdev=0.96us



Fig. 799: Histogram of release latency under GSN-EDF, with m = 24 and n = 96.











Fig. 802: Histogram of IPI latency under GSN-EDF, with m = 24 and n = 96.



Fig. 803: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 24 and n = 96.



Fig. 804: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 24 and n = 96.



Fig. 805: Histogram of release overhead under GSN-EDF-DI, with m = 24 and n = 96.

min=0.30us max=5.88us avg=1.11us median=0.73us stdev=0.95us 0.1 0.01 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.00 1.00 1.80 2.60 3.40 4.20 5.00 5.80

overhead in microseconds (bin size = 0.40us) Fig. 806: Histogram of release latency under GSN-EDF-DI, with m = 24 and n = 96.







Fig. 808: Histogram of context-switching overhead under GSN-EDF-DI, with m = 24 and n = 96.





and n = 50.



Fig. 810: Histogram of scheduling overhead (a) under G-EDF-MP, with m=24 and n=96.



Fig. 811: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 24 and n = 96.



Fig. 812: Histogram of release overhead under G-EDF-MP, with m = 24 and n = 96.



Fig. 813: Histogram of release latency under G-EDF-MP, with m=24 and n=96.

min=0.14us max=3.34us avg=0.37us median=0.27us stdev=0.24us samples: total=1513657 fraction of samples 0.1 0.01 0.00 0.0001 0.10 0.50 0.90 1.30 1.70 2.10 2.50 2.90 3.30

overhead in microseconds (bin size = 0.20us) Fig. 814: Histogram of tick overhead under G-EDF-MP, with m=24 and n=96.



Fig. 815: Histogram of context-switching overhead under G-EDF-MP, with m = 24 and n = 96.











Fig. 818: Histogram of DSP handler overhead under G-EDF-MP, with m = 24 and n = 96.



Fig. 819: Histogram of scheduling overhead (a) under SD, with m = 24and n = 96.



Fig. 820: Histogram of scheduling overhead (b) under SD, with m=24and n = 96.



Fig. 821: Histogram of release overhead under SD, with m = 24 and n = 96.

min=0.31us max=122.67us avg=3.36us median=0.76us stdev=8.57us samples: total=558099 0.1 0.01



overhead in microseconds (bin size = 10.00us) Fig. 822: Histogram of release latency under SD, with m = 24 and n = 96.



Fig. 823: Histogram of tick overhead under SD, with m = 24 and n = 96.



Fig. 824: Histogram of context-switching overhead under SD, with m = 24 and n = 96.







Fig. 826: Histogram of scheduling overhead (a) under GSN-EDF, with m = 24 and n = 120.



Fig. 827: Histogram of scheduling overhead (b) under GSN-EDF, with m = 24 and n = 120.



Fig. 828: Histogram of release overhead under GSN-EDF, with m=24and n = 120.



Fig. 829: Histogram of release latency under GSN-EDF, with m = 24and n = 120.



Fig. 830: Histogram of tick overhead under GSN-EDF, with m = 24and n = 120.





min=0.75us max=33.08us avg=1.89us median=1.90us stdev=0.47us samples: total=570711 fraction of samples 0.1 0.01 0.001 0.0001 5.00 9.00 13.00 17.00 1.00 21.00 25.00 29.00 33.00 overhead in microseconds (bin size = 2.00us)







with m = 24 and n = 120.



Fig. 834: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 24 and n = 120.



Fig. 835: Histogram of release overhead under GSN-EDF-DI, with m = 24 and n = 120.



Fig. 836: Histogram of release latency under GSN-EDF-DI, with m =24 and n = 120.



Fig. 837: Histogram of tick overhead under GSN-EDF-DI, with m = 24and n = 120.



Fig. 838: Histogram of context-switching overhead under GSN-EDF-DI, with m = 24 and n = 120.









min=0.07us max=3.26us avg=0.09us median=0.07us stdev=0.06us samples: total=763057 fraction of samples 0.1 0.01 0.001 0.0001 0.10 0.50 0.90 1.30 1.70 2.10 2.50 2.90 3.30 overhead in microseconds (bin size = 0.20us)





Fig. 842: Histogram of release overhead under G-EDF-MP, with m =24 and n = 120.



Fig. 843: Histogram of release latency under G-EDF-MP, with m = 24and n = 120.



Fig. 844: Histogram of tick overhead under G-EDF-MP, with m = 24and n = 120.



Fig. 845: Histogram of context-switching overhead under G-EDF-MP, with m = 24 and n = 120.



Fig. 846: Histogram of IPI latency under G-EDF-MP, with m=24 and n = 120.



Fig. 847: Histogram of client request latency under G-EDF-MP, with m = 24 and n = 120.



Fig. 848: Histogram of DSP handler overhead under G-EDF-MP, with m = 24 and n = 120.







Fig. 850: Histogram of scheduling overhead (b) under SD, with m = 24and n = 120.



Fig. 851: Histogram of release overhead under SD, with m = 24 and n = 120.



Fig. 852: Histogram of release latency under SD, with m = 24 and n = 120.



Fig. 853: Histogram of tick overhead under SD, with m = 24 and n = 120.



Fig. 854: Histogram of context-switching overhead under SD, with m = 24 and n = 120.







Fig. 856: Histogram of scheduling overhead (a) under GSN-EDF, with m = 24 and n = 144.



Fig. 857: Histogram of scheduling overhead (b) under GSN-EDF, with m = 24 and n = 144.



Fig. 858: Histogram of release overhead under GSN-EDF, with m = 24and n = 144.



Fig. 859: Histogram of release latency under GSN-EDF, with m = 24and n = 144.



Fig. 860: Histogram of tick overhead under GSN-EDF, with m=24and n = 144.



Fig. 861: Histogram of context-switching overhead under GSN-EDF, with m = 24 and n = 144.



Fig. 862: Histogram of IPI latency under GSN-EDF, with m = 24 and n = 144.



Fig. 863: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 24 and n = 144.











Fig. 866: Histogram of release latency under GSN-EDF-DI, with m = 24 and n = 144.



Fig. 867: Histogram of tick overhead under GSN-EDF-DI, with m = 24 and n = 144.



Fig. 868: Histogram of context-switching overhead under GSN-EDF-DI, with m = 24 and n = 144.



Fig. 869: Histogram of IPI latency under GSN-EDF-DI, with m=24 and n=144.

min=0.23us max=11.48us avg=1.50us median=1.45us stdev=0.45us samples: total=910054 fraction of samples 0.1 0.01 0.00 0.0001 0.35 1.75 3.15 4.55 5.95 7.35 8.75 10.15 11.55 overhead in microseconds (bin size = 0.70us)

Fig. 870: Histogram of scheduling overhead (a) under G-EDF-MP, with m=24 and n=144.



Fig. 871: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 24 and n = 144.



Fig. 872: Histogram of release overhead under G-EDF-MP, with m = 24 and n = 144.



Fig. 873: Histogram of release latency under G-EDF-MP, with m=24 and n=144.



Fig. 874: Histogram of tick overhead under G-EDF-MP, with m = 24 and n = 144.



Fig. 875: Histogram of context-switching overhead under G-EDF-MP, with m = 24 and n = 144.



Fig. 876: Histogram of IPI latency under G-EDF-MP, with m = 24 and n = 144.



Fig. 877: Histogram of client request latency under G-EDF-MP, with m=24 and n=144.

min=0.07us max=9.24us avg=1.49us median=1.27us stdev=0.81us samples: total=839815 0.01 0.001 0.001 0.001 0.00 0.001 0.00 0.001 0.00 0.001 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00

overhead in microseconds (bin size = 0.60us) Fig. 878: Histogram of DSP handler overhead under G-EDF-MP, with m = 24 and n = 144.



Fig. 879: Histogram of scheduling overhead (a) under SD, with m = 24 and n = 144.







Fig. 881: Histogram of release overhead under SD, with m = 24 and n = 144.



Fig. 882: Histogram of release latency under SD, with m = 24 and n = 144.



Fig. 883: Histogram of tick overhead under SD, with m = 24 and n = 144.



Fig. 884: Histogram of context-switching overhead under SD, with m = 24 and n = 144.



Fig. 885: Histogram of IPI latency under SD, with m = 24 and n =144.

min=0.33us max=152.72us avg=29.24us median=21.35us stdev=25.14us samples: total=1099550 fraction of samples 0.1 0.01 0.001 0.0001 5.00 25.00 45.00 65.00 85.00 105.00 125.00 145.00

overhead in microseconds (bin size = 10.00us) Fig. 886: Histogram of scheduling overhead (a) under GSN-EDF, with













m = 24 and n = 168.



Fig. 890: Histogram of tick overhead under GSN-EDF, with m = 24 and n = 168.



Fig. 891: Histogram of context-switching overhead under GSN-EDF, with m = 24 and n = 168.



Fig. 892: Histogram of IPI latency under GSN-EDF, with m = 24 and n = 168.



Fig. 893: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 24 and n = 168.

Fig. 894: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 24 and n = 168.

min=0.94us max=98.96us avg=29.39us median=21.86us stdev=23.54us



Fig. 895: Histogram of release overhead under GSN-EDF-DI, with m = 24 and n = 168.







Fig. 897: Histogram of tick overhead under GSN-EDF-DI, with m=24 and n=168.

min=0.07us max=2.91us avg=0.08us median=0.07us stdev=0.05us



Fig. 898: Histogram of context-switching overhead under GSN-EDF-DI, with m = 24 and n = 168.



Fig. 899: Histogram of IPI latency under GSN-EDF-DI, with m = 24and n = 168.



Fig. 900: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 24 and n = 168.



Fig. 901: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 24 and n = 168.

min=0.54us max=63.37us avg=15.20us median=11.24us stdev=12.26us samples: total=96099 fraction of samples 0. 0.01 0.001 0.0001 2.00 10.00 18.00 26.00 34.00 42.00 50.00 58.00

overhead in microseconds (bin size = 4.00us) Fig. 902: Histogram of release overhead under G-EDF-MP, with m =24 and n = 168.





Fig. 903: Histogram of release latency under G-EDF-MP, with m = 24and n = 168.





min=0.27us max=6.31us avg=1.77us median=1.80us stdev=0.42us samples: total=1115945 fraction of samples 0.1 0.01 0.001 0.0001 0.20 1.00 1.80 2.60 3.40 4.20 5.00 5.80 overhead in microseconds (bin size = 0.40us)



126



Fig. 906: Histogram of IPI latency under G-EDF-MP, with m=24 and n=168.



Fig. 907: Histogram of client request latency under G-EDF-MP, with m=24 and n=168.



Fig. 908: Histogram of DSP handler overhead under G-EDF-MP, with m = 24 and n = 168.



Fig. 909: Histogram of scheduling overhead (a) under SD, with m=24 and n=168.

min=0.06us max=668.39us avg=3.64us median=0.08us stdev=7.84us 0.01 0.001 0.001 0.0001 0.0001 0.0001 0.0001 0.0001 0.0001 0.0001 0.0001 0.000 180.00 260.00 340.00 420.00 500.00 580.00 660.00 overhead in microseconds (bin size = 40.00us)

Fig. 910: Histogram of scheduling overhead (b) under SD, with m = 24 and n = 168.



Fig. 911: Histogram of release overhead under SD, with m = 24 and n = 168.







Fig. 913: Histogram of tick overhead under SD, with m = 24 and n = 168.



Fig. 914: Histogram of context-switching overhead under SD, with m = 24 and n = 168.



Fig. 915: Histogram of IPI latency under SD, with m = 24 and n =168.



Fig. 916: Histogram of scheduling overhead (a) under GSN-EDF, with m = 24 and n = 192.



Fig. 917: Histogram of scheduling overhead (b) under GSN-EDF, with m = 24 and n = 192.

min=0.72us max=138.88us avg=35.96us median=27.54us stdev=28.64us samples: total=96121 fraction of samples 0.1 0.01 0.001 0.0001 5.00 25.00 45.00 65.00 85.00 105.00 125.00





Fig. 919: Histogram of release latency under GSN-EDF, with m = 24and n = 192.





min=0.35us max=6.39us avg=1.76us median=1.78us stdev=0.39us samples: total=1285210 fraction of samples 0.1 0.01 0.001 0.0001 0.20 1.00 1.80 2.60 3.40 4.20 5.00 5.80 overhead in microseconds (bin size = 0.40us)





Fig. 922: Histogram of IPI latency under GSN-EDF, with m = 24 and n = 192.



Fig. 923: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 24 and n = 192.



Fig. 924: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 24 and n = 192.



Fig. 925: Histogram of release overhead under GSN-EDF-DI, with m = 24 and n = 192.

samples: total=96128 fraction of samples 0.1 0.01 0.001 0.0001 0.25 1.25 2.25 3.25 4.25 5.25 6.25 7.25

overhead in microseconds (bin size = 0.50us) Fig. 926: Histogram of release latency under GSN-EDF-DI, with m =24 and n = 192.













min=0.30us max=7.53us avg=1.06us median=0.72us stdev=0.91us



Fig. 930: Histogram of scheduling overhead (a) under G-EDF-MP, with m=24 and n=192.



Fig. 931: Histogram of scheduling overhead (b) under G-EDF-MP, with m=24 and n=192.



Fig. 932: Histogram of release overhead under G-EDF-MP, with m = 24 and n = 192.



Fig. 933: Histogram of release latency under G-EDF-MP, with m=24 and n=192.

 $\underset{t_{i}}{\text{min=0.13us max=2.44us avg=0.36us media=0.27us stdev=0.23us}} \\ \underset{t_{i}}{\text{min=0.13us max=2.44us avg=0.36us media=0.27us stdev=0.27us stdev=0.27us$



min=0.27us max=6.13us avg=1.76us median=1.79us stdev=0.43us



Fig. 935: Histogram of context-switching overhead under G-EDF-MP, with m = 24 and n = 192.











Fig. 938: Histogram of DSP handler overhead under G-EDF-MP, with m=24 and n=192.



Fig. 939: Histogram of scheduling overhead (a) under SD, with m = 24 and n = 192.



Fig. 940: Histogram of scheduling overhead (b) under SD, with m=24 and n=192.



Fig. 941: Histogram of release overhead under SD, with m=24 and n=192.

samples: total=1009268 0.01 0.001 7.50 37.50 67.50 97.50 127.50 157.50 187.50

overhead in microseconds (bin size = 15.00us) Fig. 942: Histogram of release latency under SD, with m = 24 and n = 192.



Fig. 943: Histogram of tick overhead under SD, with m = 24 and n = 192.







Fig. 945: Histogram of IPI latency under SD, with m = 24 and n = 192.

min=0.28us max=197.91us avg=8.95us median=0.71us stdev=19.03us



Fig. 946: Histogram of scheduling overhead (a) under GSN-EDF, with m=24 and n=216.



Fig. 947: Histogram of scheduling overhead (b) under GSN-EDF, with m = 24 and n = 216.



Fig. 948: Histogram of release overhead under GSN-EDF, with m = 24 and n = 216.



Fig. 949: Histogram of release latency under GSN-EDF, with m=24 and n=216.



Fig. 950: Histogram of tick overhead under GSN-EDF, with m = 24 and n = 216.





Fig. 951: Histogram of context-switching overhead under GSN-EDF, with m = 24 and n = 216.









8us stdev=28.81us min=0.7



Fig. 954: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 24 and n = 216.



Fig. 955: Histogram of release overhead under GSN-EDF-DI, with m = 24 and n = 216.



Fig. 956: Histogram of release latency under GSN-EDF-DI, with m = 24 and n = 216.



Fig. 957: Histogram of tick overhead under GSN-EDF-DI, with m=24 and n=216.



Fig. 958: Histogram of context-switching overhead under GSN-EDF-DI, with m = 24 and n = 216.







Fig. 960: Histogram of scheduling overhead (a) under G-EDF-MP, with m=24 and n=216.





samt Samt



Fig. 962: Histogram of release overhead under G-EDF-MP, with m = 24 and n = 216.



Fig. 963: Histogram of release latency under G-EDF-MP, with m = 24 and n = 216.



Fig. 964: Histogram of tick overhead under G-EDF-MP, with m = 24 and n = 216.



Fig. 965: Histogram of context-switching overhead under G-EDF-MP, with m=24 and n=216.



Fig. 966: Histogram of IPI latency under G-EDF-MP, with m=24 and n=216.



Fig. 967: Histogram of client request latency under G-EDF-MP, with m = 24 and n = 216.









m = 24 and r



Fig. 970: Histogram of scheduling overhead (b) under SD, with m = 24and n = 216.



Fig. 971: Histogram of release overhead under SD, with m = 24 and n = 216.



Fig. 972: Histogram of release latency under SD, with m = 24 and n = 216.



Fig. 973: Histogram of tick overhead under SD, with m = 24 and n = 216.

min=0.25us max=8.44us avg=1.60us median=1.62us stdev=0.48us samples: total=1484351 fraction of samples 0.1 0.01 0.00 0.0001 0.25 1.25 2.25 3.25 4.25 5.25 6.25 7.25 8.25 overhead in microseconds (bin size = 0.50us)









Fig. 976: Histogram of scheduling overhead (a) under GSN-EDF, with m = 24 and n = 240.





135



Fig. 978: Histogram of release overhead under GSN-EDF, with m=24 and n=240.



Fig. 979: Histogram of release latency under GSN-EDF, with m = 24 and n = 240.



Fig. 980: Histogram of tick overhead under GSN-EDF, with m = 24 and n = 240.



Fig. 981: Histogram of context-switching overhead under GSN-EDF, with m=24 and n=240.







Fig. 983: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 24 and n = 240.









us min=0.07us ma



Fig. 986: Histogram of release latency under GSN-EDF-DI, with m =24 and n = 240.



Fig. 987: Histogram of tick overhead under GSN-EDF-DI, with m=24and n = 240.



Fig. 988: Histogram of context-switching overhead under GSN-EDF-DI, with m = 24 and n = 240.



Fig. 989: Histogram of IPI latency under GSN-EDF-DI, with m = 24and n = 240.

min=0.20us max=13.15us avg=1.44us median=1.40us stdev=0.42us samples: total=1627296 fraction of samples 0.1 0.01 0.001 0.0001 0.40 2.00 3.60 5.20 6.80 8.40 10.00 11.60 13.20 overhead in microseconds (bin size = 0.80us)

Fig. 990: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 24 and n = 240.



Fig. 991: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 24 and n = 240.





min=0.26us max=8.32us avg=1.01us median=0.68us stdev=0.89us samples: total=96900 fraction of samples 0.1 0.01 0.001 0.0001 0.25 1.25 2 25 3 25 4.25 5.25 6.25 7.25 8.25 overhead in microseconds (bin size = 0.50us)





Fig. 994: Histogram of tick overhead under G-EDF-MP, with m = 24 and n = 240.



Fig. 995: Histogram of context-switching overhead under G-EDF-MP, with m=24 and n=240.



Fig. 996: Histogram of IPI latency under G-EDF-MP, with m = 24 and n = 240.



Fig. 997: Histogram of client request latency under G-EDF-MP, with m=24 and n=240.

min=0.07us max=9.93us avg=1.78us median=1.38us stdev=1.05us samples: total=1429798 fraction of samples 0.1 0.01 0.00 0.0001 0.30 1.50 2.70 3.90 5.10 7.50 8.70 9.90 overhead in microseconds (bin size = 0.60us)

Fig. 998: Histogram of DSP handler overhead under G-EDF-MP, with m = 24 and n = 240.











Fig. 1001: Histogram of release overhead under SD, with m = 24 and n = 240.



Fig. 1002: Histogram of release latency under SD, with m=24 and n=240.



Fig. 1003: Histogram of tick overhead under SD, with m = 24 and n = 240.



Fig. 1004: Histogram of context-switching overhead under SD, with m = 24 and n = 240.



Fig. 1005: Histogram of IPI latency under SD, with m = 24 and n = 240.

8 0.01 0.001 0.001 7.50 37.50 67.50 97.50 127.50 157.50 187.50

overhead in microseconds (bin size = 15.00us) Fig. 1006: Histogram of scheduling overhead (a) under GSN-EDF, with m = 32 and n = 32.



Fig. 1007: Histogram of scheduling overhead (b) under GSN-EDF, with m = 32 and n = 32.



Fig. 1008: Histogram of release overhead under GSN-EDF, with m = 32 and n = 32.



Fig. 1009: Histogram of release latency under GSN-EDF, with m=32 and n=32.

min=0.28us max=189.65us avg=8.38us median=4.24us stdev=10.22us



Fig. 1010: Histogram of tick overhead under GSN-EDF, with m = 32 and n = 32.



Fig. 1011: Histogram of context-switching overhead under GSN-EDF, with m = 32 and n = 32.



Fig. 1012: Histogram of IPI latency under GSN-EDF, with m = 32 and n = 32.



Fig. 1013: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 32 and n = 32.

min=0.07us max=2.12us avg=0.09us median=0.07us stdev=0.07us



Fig. 1014: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 32 and n = 32.



Fig. 1015: Histogram of release overhead under GSN-EDF-DI, with m = 32 and n = 32.





min=0.13us max=4.75us avg=0.38us median=0.20us stdev=0.31us samples: total=2521881 fraction of samples 0.1 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 4.35 overhead in microseconds (bin size = 0.30us)





Fig. 1018: Histogram of context-switching overhead under GSN-EDF-DI, with m = 32 and n = 32.



Fig. 1019: Histogram of IPI latency under GSN-EDF-DI, with m = 32and n = 32.



Fig. 1020: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 32 and n = 32.



Fig. 1021: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 32 and n = 32.

samples: total=72422 fraction of samples 0.1 0.01 0.001 0.0001 1.00 5.00 9.00 13.00 17.00 21.00 25.00 29.00 33.00 overhead in microseconds (bin size = 2.00us)

Fig. 1022: Histogram of release overhead under G-EDF-MP, with m =32 and n = 32.

min=0.31us max=7.95us avg=1.45us median=0.79us stdev=1.31us



Fig. 1023: Histogram of release latency under G-EDF-MP, with m = 32and n = 32.









min=0.45us max=33.05us avg=4.37us median=3.08us stdev=3.71us



Fig. 1026: Histogram of IPI latency under G-EDF-MP, with m = 32 and n = 32.



Fig. 1027: Histogram of client request latency under G-EDF-MP, with m = 32 and n = 32.



Fig. 1028: Histogram of DSP handler overhead under G-EDF-MP, with m = 32 and n = 32.



Fig. 1029: Histogram of scheduling overhead (a) under SD, with m = 32 and n = 32.

min=0.06us max=114.58us avg=0.54us median=0.07us stdev=1.97us 0.01 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.0010.



min=0.28us max=54.29us avg=5.72us median=3.51us stdev=5.70us 0.01 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.00

Fig. 1031: Histogram of release overhead under SD, with m = 32 and n = 32.







Fig. 1033: Histogram of tick overhead under SD, with m = 32 and n = 32.



Fig. 1034: Histogram of context-switching overhead under SD, with m = 32 and n = 32.



Fig. 1035: Histogram of IPI latency under SD, with m = 32 and n = 32.



Fig. 1036: Histogram of scheduling overhead (a) under GSN-EDF, with m=32 and n=64.



Fig. 1037: Histogram of scheduling overhead (b) under GSN-EDF, with m=32 and n=64.

min=0.42us max=167.10us avg=15.60us median=9.47us stdev=16.76us



Fig. 1038: Histogram of release overhead under GSN-EDF, with m = 32 and n = 64.

min=0.29us max=9.27us avg=1.39us median=0.80us stdev=1.27us



Fig. 1039: Histogram of release latency under GSN-EDF, with m=32 and n=64.











Fig. 1042: Histogram of IPI latency under GSN-EDF, with m = 32 and n = 64.



Fig. 1043: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m=32 and n=64.



Fig. 1044: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 32 and n = 64.



Fig. 1045: Histogram of release overhead under GSN-EDF-DI, with m=32 and n=64.

min=0.30us max=7.53us avg=1.45us median=0.81us stdev=1.29us samples: total=88489 fraction of samples 0.1 0.01 0.001 0.0001 0.25 1.25 2.25 3.25 4.25 5.25 6.25 7.25 overhead in microseconds (bin size = 0.50us)

Fig. 1046: Histogram of release latency under GSN-EDF-DI, with m = 32 and n = 64.







Fig. 1048: Histogram of context-switching overhead under GSN-EDF-DI, with m = 32 and n = 64.



Fig. 1049: Histogram of IPI latency under GSN-EDF-DI, with m=32 and n=64.


Fig. 1050: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 32 and n = 64.



Fig. 1051: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 32 and n = 64.



Fig. 1052: Histogram of release overhead under G-EDF-MP, with m =32 and n = 64.



Fig. 1053: Histogram of release latency under G-EDF-MP, with m = 32and n = 64.

min=0.13us max=3.96us avg=0.41us median=0.31us stdev=0.31us samples: total=2418270 fraction of samples 0.1 0.01 0.001 0.0001 3.75 0.15 0.75 1.35 1.95 2.55 3.15 overhead in microseconds (bin size = 0.30us)





2.60 overhead in microseconds (bin size = 0.40us) Fig. 1055: Histogram of context-switching overhead under G-EDF-MP, with m = 32 and n = 64.

3.40

4.20

5.00

5.80

0.0001

0.20

1.00

1.80











Fig. 1058: Histogram of DSP handler overhead under G-EDF-MP, with m = 32 and n = 64.



Fig. 1059: Histogram of scheduling overhead (a) under SD, with m =32 and n = 64.



Fig. 1060: Histogram of scheduling overhead (b) under SD, with m =32 and n = 64.



Fig. 1061: Histogram of release overhead under SD, with m = 32 and n = 64.

min=0.30us max=135.64us avg=1.61us median=0.68us stdev=5.52us samples: total=379737 fraction of samples 0.1 0.01 0.001 0.0001 125.00 5.00 25.00 45.00 65.00 85.00 105.00 overhead in microseconds (bin size = 10.00us)

Fig. 1062: Histogram of release latency under SD, with m = 32 and n = 64.



Fig. 1063: Histogram of tick overhead under SD, with m = 32 and n = 64.











Fig. 1066: Histogram of scheduling overhead (a) under GSN-EDF, with m = 32 and n = 96.



Fig. 1067: Histogram of scheduling overhead (b) under GSN-EDF, with m = 32 and n = 96.



Fig. 1068: Histogram of release overhead under GSN-EDF, with m =32 and n = 96.



Fig. 1069: Histogram of release latency under GSN-EDF, with m = 32and n = 96.

min=0.13us max=4.77us avg=0.42us median=0.28us stdev=0.31us samples: total=2500459 fraction of samples 0.1 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 4.35 overhead in microseconds (bin size = 0.30us)

Fig. 1070: Histogram of tick overhead under GSN-EDF, with m = 32and n = 96.





Fig. 1071: Histogram of context-switching overhead under GSN-EDF, with m = 32 and n = 96.











Fig. 1074: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 32 and n = 96.



Fig. 1075: Histogram of release overhead under GSN-EDF-DI, with m = 32 and n = 96.



Fig. 1076: Histogram of release latency under GSN-EDF-DI, with m =32 and n = 96.



Fig. 1077: Histogram of tick overhead under GSN-EDF-DI, with m =32 and n = 96.



Fig. 1078: Histogram of context-switching overhead under GSN-EDF-DI, with m = 32 and n = 96.



overhead in microseconds (bin size = 0.40us) Fig. 1079: Histogram of IPI latency under GSN-EDF-DI, with m = 32and n = 96.



Fig. 1080: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 32 and n = 96.







Fig. 1082: Histogram of release overhead under G-EDF-MP, with m = 32 and n = 96.



Fig. 1083: Histogram of release latency under G-EDF-MP, with m = 32 and n = 96.



Fig. 1084: Histogram of tick overhead under G-EDF-MP, with m = 32 and n = 96.



Fig. 1085: Histogram of context-switching overhead under G-EDF-MP, with m = 32 and n = 96.



Fig. 1086: Histogram of IPI latency under G-EDF-MP, with m = 32 and n = 96.











Fig. 1089: Histogram of scheduling overhead (a) under SD, with m = 32 and n = 96.



Fig. 1090: Histogram of scheduling overhead (b) under SD, with m =32 and n = 96.



Fig. 1091: Histogram of release overhead under SD, with m = 32 and n = 96.



Fig. 1092: Histogram of release latency under SD, with m = 32 and n = 96.



Fig. 1093: Histogram of tick overhead under SD, with m = 32 and n = 96.

min=0.26us max=12.26us avg=1.56us median=1.57us stdev=0.60us samples: total=736827 fraction of samples 0. 0.01 0.00 0.0001 0.35 1.75 3.15 4.55 5.95 7.35 8.75 10.15 11.55 overhead in microseconds (bin size = 0.70us)









Fig. 1096: Histogram of scheduling overhead (a) under GSN-EDF, with m = 32 and n = 128.







Fig. 1098: Histogram of release overhead under GSN-EDF, with m = 32 and n = 128.



Fig. 1099: Histogram of release latency under GSN-EDF, with m = 32 and n = 128.



Fig. 1100: Histogram of tick overhead under GSN-EDF, with m = 32 and n = 128.



Fig. 1101: Histogram of context-switching overhead under GSN-EDF, with m = 32 and n = 128.

min=0.72us max=20.34us avg=2.02us median=1.98us stdev=0.55us 0.01 0.001 0.001 0.001 1.00 5.00 9.00 13.00 17.00 21.00overhead in microseconds (bin size = 2.00us)



min=0.34us max=185.13us avg=34.00us median=25.66us stdev=29.05us



Fig. 1103: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 32 and n = 128.



Fig. 1104: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 32 and n = 128.







Fig. 1106: Histogram of release latency under GSN-EDF-DI, with m =32 and n = 128.



Fig. 1107: Histogram of tick overhead under GSN-EDF-DI, with m =32 and n = 128.



Fig. 1108: Histogram of context-switching overhead under GSN-EDF-DI, with m = 32 and n = 128.



Fig. 1109: Histogram of IPI latency under GSN-EDF-DI, with m = 32and n = 128.

min=0.21us max=12.78us avg=1.67us median=1.62us stdev=0.46us samples: total=776025 fraction of samples 0.1 0.01 0.001 0.0001 0.40 2.00 3.60 5.20 6.80 8.40 10.00 11.60 overhead in microseconds (bin size = 0.80us)

Fig. 1110: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 32 and n = 128.









min=0.31us max=8.95us avg=1.26us median=0.73us stdev=1.21us samples: total=92960 fraction of samples 0.1 0.01 0.001 0.0001 0.25 1 25 2 25 3 25 4.25 5.25 6.25 7.25 8.25 overhead in microseconds (bin size = 0.50us)





Fig. 1114: Histogram of tick overhead under G-EDF-MP, with m = 32 and n = 128.



Fig. 1115: Histogram of context-switching overhead under G-EDF-MP, with m = 32 and n = 128.



Fig. 1116: Histogram of IPI latency under G-EDF-MP, with m = 32 and n = 128.



Fig. 1117: Histogram of client request latency under G-EDF-MP, with m=32 and n=128.

min=0.07us max=10.01us avg=1.40us median=1.25us stdev=0.68us samples: total=751878 fraction of samples 0. 0.01 0.00 0.0001 0.30 1.50 2.70 3.90 5.10 7.50 8.70 9.90 overhead in microseconds (bin size = 0.60us)

Fig. 1118: Histogram of DSP handler overhead under G-EDF-MP, with m = 32 and n = 128.











Fig. 1121: Histogram of release overhead under SD, with m = 32 and n = 128.



Fig. 1122: Histogram of release latency under SD, with m = 32 and n = 128.



Fig. 1123: Histogram of tick overhead under SD, with m = 32 and n = 128.



Fig. 1124: Histogram of context-switching overhead under SD, with m = 32 and n = 128.



Fig. 1125: Histogram of IPI latency under SD, with m = 32 and n = 128.

min=0.33us max=236.33us avg=44.68us median=32.03us stdev=40.42us



Fig. 1126: Histogram of scheduling overhead (a) under GSN-EDF, with m=32 and n=160.



Fig. 1127: Histogram of scheduling overhead (b) under GSN-EDF, with m=32 and n=160.











Fig. 1130: Histogram of tick overhead under GSN-EDF, with m=32 and n=160.



Fig. 1131: Histogram of context-switching overhead under GSN-EDF, with m = 32 and n = 160.



Fig. 1132: Histogram of IPI latency under GSN-EDF, with m = 32 and n = 160.



Fig. 1133: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m=32 and n=160.

source in the second se

Fig. 1134: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 32 and n = 160.



overhead in microseconds (bin size = 10.00us) Fig. 1135: Histogram of release overhead under GSN-EDF-DI, with m = 32 and n = 160.



Fig. 1136: Histogram of release latency under GSN-EDF-DI, with m = 32 and n = 160.



Fig. 1137: Histogram of tick overhead under GSN-EDF-DI, with m=32 and n=160.

min=0.07us max=4.78us avg=0.08us median=0.07us stdev=0.07us



Fig. 1138: Histogram of context-switching overhead under GSN-EDF-DI, with m = 32 and n = 160.



Fig. 1139: Histogram of IPI latency under GSN-EDF-DI, with m = 32 and n = 160.



Fig. 1140: Histogram of scheduling overhead (a) under G-EDF-MP, with m=32 and n=160.



Fig. 1141: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 32 and n = 160.

min=0.53us max=82.34us avg=17.10us median=12.00us stdev=15.34us



Fig. 1142: Histogram of release overhead under G-EDF-MP, with m = 32 and n = 160.

min=0.30us max=8.99us avg=1.20us median=0.69us stdev=1.19us



Fig. 1143: Histogram of release latency under G-EDF-MP, with m = 32 and n = 160.



Fig. 1144: Histogram of tick overhead under G-EDF-MP, with m = 32 and n = 160.



Fig. 1145: Histogram of context-switching overhead under G-EDF-MP, with m=32 and n=160.



Fig. 1146: Histogram of IPI latency under G-EDF-MP, with m = 32and n = 160.



Fig. 1147: Histogram of client request latency under G-EDF-MP, with m = 32 and n = 160.



Fig. 1148: Histogram of DSP handler overhead under G-EDF-MP, with m = 32 and n = 160.



Fig. 1149: Histogram of scheduling overhead (a) under SD, with m =32 and n = 160.

min=0.06us max=445.35us avg=4.98us median=0.08us stdev=11.92us samples: total=1176626 fraction of samples 0.1 0.01 0.00 0.0001 12.50 62.50 112.50 162.50 212.50 262.50 312.50 362.50 412.50

overhead in microseconds (bin size = 25.00us) Fig. 1150: Histogram of scheduling overhead (b) under SD, with m =32 and n = 160.



Fig. 1151: Histogram of release overhead under SD, with m = 32 and n = 160.











Fig. 1154: Histogram of context-switching overhead under SD, with m = 32 and n = 160.



Fig. 1155: Histogram of IPI latency under SD, with m = 32 and n = 160.



Fig. 1156: Histogram of scheduling overhead (a) under GSN-EDF, with m = 32 and n = 192.



Fig. 1157: Histogram of scheduling overhead (b) under GSN-EDF, with m=32 and n=192.

min=0.64us max=193.01us avg=48.01us median=33.07us stdev=42.05us



Fig. 1158: Histogram of release overhead under GSN-EDF, with m = 32 and n = 192.



Fig. 1159: Histogram of release latency under GSN-EDF, with m = 32 and n = 192.



Fig. 1160: Histogram of tick overhead under GSN-EDF, with m = 32 and n = 192.







Fig. 1162: Histogram of IPI latency under GSN-EDF, with m = 32 and n = 192.



Fig. 1163: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m=32 and n=192.



Fig. 1164: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 32 and n = 192.



Fig. 1165: Histogram of release overhead under GSN-EDF-DI, with m = 32 and n = 192.

 $\begin{array}{c} \text{min=0.33us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \\ \begin{array}{c} \text{min=0.33us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \\ \begin{array}{c} \text{min=0.33us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.33us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.33us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.33us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.33us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.33us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.33us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.33us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.33us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.33us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.33us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.33us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.33us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.33us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.33us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.31us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.31us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.31us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.31us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.31us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.31us max=6.87us avg=1.39us median=0.81us stdev=1.25us} \\ \end{array} \\ \begin{array}{c} \text{min=0.31us max=6.87us avg=1.25us} \\ \end{array} \\ \end{array} \\ \begin{array}{c} \text{min=0.31us max=0.39us max=6.87us max=0.39us} \\ \end{array} \\ \begin{array}{c} \text{min=0.31us max=0.39us} \\ \end{array} \\ \end{array} \\ \begin{array}{c} \text{min=0.31us max=0.39us} \\ \end{array} \\ \end{array} \\ \end{array} \\ \begin{array}{c} \text{min=0.31us max=0.39us} \\ \end{array} \\ \end{array} \\ \end{array} \\ \end{array} \\ \begin{array}{c} \text{min=0.31us max=0.39us} \\ \end{array} \\ \end{array} \\ \end{array} \\ \end{array} \\ \end{array} \\ \end{array}$ \\ \begin{array}{c} \text{mi=0.31us max=0.39us} \\ \end{array} \\ \end{array} \\ \end{array} \\ \end{array} \\













Fig. 1169: Histogram of IPI latency under GSN-EDF-DI, with m = 32 and n = 192.



Fig. 1170: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 32 and n = 192.



Fig. 1171: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 32 and n = 192.



Fig. 1172: Histogram of release overhead under G-EDF-MP, with m = 32 and n = 192.



Fig. 1173: Histogram of release latency under G-EDF-MP, with m=32 and n=192.

min=0.13us max=4.29us avg=0.40us median=0.27us stdev=0.29us samples: total=1926891 fraction of samples 0.1 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 4.35 overhead in microseconds (bin size = 0.30us)

Fig. 1174: Histogram of tick overhead under G-EDF-MP, with m = 32 and n = 192.



Fig. 1175: Histogram of context-switching overhead under G-EDF-MP, with m = 32 and n = 192.







Fig. 1177: Histogram of client request latency under G-EDF-MP, with m=32 and n=192.



Fig. 1178: Histogram of DSP handler overhead under G-EDF-MP, with m = 32 and n = 192.



Fig. 1179: Histogram of scheduling overhead (a) under SD, with m = 32 and n = 192.



Fig. 1180: Histogram of scheduling overhead (b) under SD, with m = 32 and n = 192.



Fig. 1181: Histogram of release overhead under SD, with m = 32 and n = 192.

min=0.27us max=277.12us avg=9.78us median=0.75us stdev=25.33us



Fig. 1182: Histogram of release latency under SD, with m = 32 and n = 192.



Fig. 1183: Histogram of tick overhead under SD, with m = 32 and n = 192.







Fig. 1185: Histogram of IPI latency under SD, with m = 32 and n = 192.



Fig. 1186: Histogram of scheduling overhead (a) under GSN-EDF, with m = 32 and n = 224.



Fig. 1187: Histogram of scheduling overhead (b) under GSN-EDF, with m = 32 and n = 224.



Fig. 1188: Histogram of release overhead under GSN-EDF, with m =32 and n = 224.



Fig. 1189: Histogram of release latency under GSN-EDF, with m = 32and n = 224.

min=0.14us max=4.48us avg=0.45us median=0.40us stdev=0.28us samples: total=1935107 fraction of samples 0.1 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 4.35 overhead in microseconds (bin size = 0.30us)

Fig. 1190: Histogram of tick overhead under GSN-EDF, with m = 32and n = 224.



Fig. 1191: Histogram of context-switching overhead under GSN-EDF, with m = 32 and n = 224.











Fig. 1194: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 32 and n = 224.



Fig. 1195: Histogram of release overhead under GSN-EDF-DI, with m = 32 and n = 224.



Fig. 1196: Histogram of release latency under GSN-EDF-DI, with m =32 and n = 224.



Fig. 1197: Histogram of tick overhead under GSN-EDF-DI, with m =32 and n = 224.



Fig. 1198: Histogram of context-switching overhead under GSN-EDF-DI, with m = 32 and n = 224.



Fig. 1199: Histogram of IPI latency under GSN-EDF-DI, with m = 32and n = 224.



Fig. 1200: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 32 and n = 224.







Fig. 1202: Histogram of release overhead under G-EDF-MP, with m = 32 and n = 224.



Fig. 1203: Histogram of release latency under G-EDF-MP, with m=32 and n=224.



Fig. 1204: Histogram of tick overhead under G-EDF-MP, with m = 32 and n = 224.



Fig. 1205: Histogram of context-switching overhead under G-EDF-MP, with m = 32 and n = 224.



Fig. 1206: Histogram of IPI latency under G-EDF-MP, with m = 32 and n = 224.



Fig. 1207: Histogram of client request latency under G-EDF-MP, with m = 32 and n = 224.











Fig. 1210: Histogram of scheduling overhead (b) under SD, with m = 32 and n = 224.



Fig. 1211: Histogram of release overhead under SD, with m = 32 and n = 224.



Fig. 1212: Histogram of release latency under SD, with m = 32 and n = 224.



Fig. 1213: Histogram of tick overhead under SD, with m = 32 and n = 224.

Fig. 1214: Histogram of context-switching overhead under SD, with m = 32 and n = 224.

overhead in microseconds (bin size = 0.80us)







Fig. 1216: Histogram of scheduling overhead (a) under GSN-EDF, with m=32 and n=256.



Fig. 1217: Histogram of scheduling overhead (b) under GSN-EDF, with $m=32 \mbox{ and } n=256.$

min=0.23us max=13.75us avg=1.72us median=1.68us stdev=0.65us



Fig. 1218: Histogram of release overhead under GSN-EDF, with m = 32 and n = 256.



Fig. 1219: Histogram of release latency under GSN-EDF, with m=32 and n=256.



Fig. 1220: Histogram of tick overhead under GSN-EDF, with m=32 and n=256.



Fig. 1221: Histogram of context-switching overhead under GSN-EDF, with m = 32 and n = 256.

min=0.70us max=109.73us avg=1.97us median=1.94us stdev=0.58us 0.01 0.001 0.001 0.001 5.00 25.00 45.00 65.00 85.00 105.00 overhead in microseconds (bin size = 10.00us)





Fig. 1223: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 32 and n = 256.



Fig. 1224: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 32 and n = 256.



Fig. 1225: Histogram of release overhead under GSN-EDF-DI, with $m=32 \ {\rm and} \ n=256.$



Fig. 1226: Histogram of release latency under GSN-EDF-DI, with m = 32 and n = 256.



Fig. 1227: Histogram of tick overhead under GSN-EDF-DI, with m = 32 and n = 256.



Fig. 1228: Histogram of context-switching overhead under GSN-EDF-DI, with m = 32 and n = 256.



Fig. 1229: Histogram of IPI latency under GSN-EDF-DI, with m=32 and n=256.

min=0.21us max=12.77us avg=1.55us median=1.51us stdev=0.42us samples: total=11777584 fraction of samples 0.1 0.01 0.001 0.0001 0.40 2.00 3.60 5.20 6.80 8.40 10.00 11.60 overhead in microseconds (bin size = 0.80us)

Fig. 1230: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 32 and n = 256.







Fig. 1232: Histogram of release overhead under G-EDF-MP, with m=32 and n=256.



Fig. 1233: Histogram of release latency under G-EDF-MP, with m=32 and n=256.



Fig. 1234: Histogram of tick overhead under G-EDF-MP, with m=32 and n=256.



Fig. 1235: Histogram of context-switching overhead under G-EDF-MP, with m=32 and n=256.



Fig. 1236: Histogram of IPI latency under G-EDF-MP, with m = 32 and n = 256.



Fig. 1237: Histogram of client request latency under G-EDF-MP, with m=32 and n=256.

min=0.07us max=10.07us avg=1.72us median=1.31us stdev=1.02us samples: total=1553831 fraction of samples 0.1 0.01 0.00 0.0001 0.30 1.50 2.70 3.90 5.10 7.50 8.70 9.90 overhead in microseconds (bin size = 0.60us)

Fig. 1238: Histogram of DSP handler overhead under G-EDF-MP, with m = 32 and n = 256.











Fig. 1241: Histogram of release overhead under SD, with m=32 and n=256.



Fig. 1242: Histogram of release latency under SD, with m = 32 and n = 256.



Fig. 1243: Histogram of tick overhead under SD, with m = 32 and n = 256.



Fig. 1244: Histogram of context-switching overhead under SD, with m = 32 and n = 256.



Fig. 1245: Histogram of IPI latency under SD, with m = 32 and n = 256.

samples: total=1924154 fraction of samples 0.1 0.01 0.00 0.0001 10.00 50.00 90.00 130.00 170.00 210.00 250.00 290.00 330.00





Fig. 1247: Histogram of scheduling overhead (b) under GSN-EDF, with m = 32 and n = 288.





min=0.30us max=14.26us avg=1.57us median=0.93us stdev=1.33us samples: total=97353 fraction of samples 0.1 0.01 0.001 0.0001 0.40 2.00 3.60 5.20 6.80 8.40 10.00 11.60 13.20 overhead in microseconds (bin size = 0.80us)



min=0.33us max=321.37us avg=57.66us median=46.96us stdev=46.91us



Fig. 1250: Histogram of tick overhead under GSN-EDF, with m = 32and n = 288.



Fig. 1251: Histogram of context-switching overhead under GSN-EDF, with m = 32 and n = 288.



Fig. 1252: Histogram of IPI latency under GSN-EDF, with m = 32 and n = 288.



Fig. 1253: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 32 and n = 288.

min=0.07us max=4.24us avg=0.08us median=0.07us stdev=0.06us samples: total=1940878 fraction of samples 0.1 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 4.35 overhead in microseconds (bin size = 0.30us)

Fig. 1254: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 32 and n = 288.



Fig. 1255: Histogram of release overhead under GSN-EDF-DI, with m = 32 and n = 288.





min=0.13us max=5.02us avg=0.45us median=0.41us stdev=0.26us samples: total=1917057 fraction of samples 0.1 0.01 0.001 0.0001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 4.35 4.95 overhead in microseconds (bin size = 0.30us)





Fig. 1258: Histogram of context-switching overhead under GSN-EDF-DI, with m = 32 and n = 288.



Fig. 1259: Histogram of IPI latency under GSN-EDF-DI, with m = 32 and n = 288.



Fig. 1260: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 32 and n = 288.



Fig. 1261: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 32 and n = 288.

min=0.58us max=84.25us avg=27.69us median=23.01us stdev=19.75us samples: total=97352 fraction of samples 0. 0.01 0.00 0.0001 2.50 12.50 22.50 32.50 42.50 52.50 62.50 72.50 82.50



min=0.29us max=8.93us avg=1.16us median=0.70us stdev=1.14us



Fig. 1263: Histogram of release latency under G-EDF-MP, with m = 32 and n = 288.



Fig. 1264: Histogram of tick overhead under G-EDF-MP, with m = 32 and n = 288.



Fig. 1265: Histogram of context-switching overhead under G-EDF-MP, with m=32 and n=288.



Fig. 1266: Histogram of IPI latency under G-EDF-MP, with m = 32and n = 288.



Fig. 1267: Histogram of client request latency under G-EDF-MP, with m = 32 and n = 288.



Fig. 1268: Histogram of DSP handler overhead under G-EDF-MP, with m = 32 and n = 288.



Fig. 1269: Histogram of scheduling overhead (a) under SD, with m =32 and n = 288.

min=0.06us max=627.90us avg=6.41us median=0.09us stdev=15.85us samples: total=1968283 fraction of samples 0.1 0.01 0.001 0.0001 20.00 100.00 180.00 260.00 340.00 420.00 500.00 580.00 overhead in microseconds (bin size = 40.00us)



min=0.19us max=389.95us avg=29.03us median=20.64us stdev=29.85us















Fig. 1274: Histogram of context-switching overhead under SD, with m = 32 and n = 288.



Fig. 1275: Histogram of IPI latency under SD, with m = 32 and n = 288.



Fig. 1276: Histogram of scheduling overhead (a) under GSN-EDF, with m = 32 and n = 320.



Fig. 1277: Histogram of scheduling overhead (b) under GSN-EDF, with m=32 and n=320.

min=0.82us max=209.19us avg=73.94us median=63.72us stdev=52.19us



Fig. 1278: Histogram of release overhead under GSN-EDF, with m = 32 and n = 320.







Fig. 1280: Histogram of tick overhead under GSN-EDF, with m = 32 and n = 320.



Fig. 1281: Histogram of context-switching overhead under GSN-EDF, with m=32 and n=320.



Fig. 1282: Histogram of IPI latency under GSN-EDF, with m = 32 and n = 320.



Fig. 1283: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 32 and n = 320.



Fig. 1284: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 32 and n = 320.



Fig. 1285: Histogram of release overhead under GSN-EDF-DI, with m=32 and n=320.

min=0.32us max=6.83us avg=1.36us median=0.79us stdev=1.24us samples: total=97034 fraction of samples 0.1 0.01 0.001 0.0001 0.20 1.00 1.80 2.60 3.40 4.20 5.00 5.80 6.60



min=0.13us max=3.84us avg=0.46us median=0.41us stdev=0.26us samples: total=2392336 0.01 0.001 0.001 0.15 0.75 1.35 1.95 2.55 3.15 3.75 overhead in microseconds (bin size = 0.30us)

Fig. 1287: Histogram of tick overhead under GSN-EDF-DI, with m = 32 and n = 320.







Fig. 1289: Histogram of IPI latency under GSN-EDF-DI, with m=32 and n=320.



Fig. 1290: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 32 and n = 320.



Fig. 1291: Histogram of scheduling overhead (b) under G-EDF-MP, with m=32 and n=320.



Fig. 1292: Histogram of release overhead under G-EDF-MP, with m = 32 and n = 320.



Fig. 1293: Histogram of release latency under G-EDF-MP, with m=32 and n=320.

min=0.13us max=3.98us avg=0.38us median=0.26us stdev=0.27us 0.01 0.001 0.001 0.15 0.75 1.35 1.95 2.55 3.15 3.75



min=0.30us max=6.31us avg=1.75us median=1.77us stdev=0.44us



Fig. 1295: Histogram of context-switching overhead under G-EDF-MP, with m = 32 and n = 320.











Fig. 1298: Histogram of DSP handler overhead under G-EDF-MP, with m=32 and n=320.



Fig. 1299: Histogram of scheduling overhead (a) under SD, with m = 32 and n = 320.



Fig. 1300: Histogram of scheduling overhead (b) under SD, with m = 32 and n = 320.



Fig. 1301: Histogram of release overhead under SD, with m = 32 and n = 320.

min=0.28us max=508.28us avg=24.23us median=0.86us stdev=51.04us



Fig. 1302: Histogram of release latency under SD, with m = 32 and n = 320.







Fig. 1304: Histogram of context-switching overhead under SD, with m = 32 and n = 320.







Fig. 1306: Histogram of scheduling overhead (a) under GSN-EDF, with m = 48 and n = 48.



Fig. 1307: Histogram of scheduling overhead (b) under GSN-EDF, with m = 48 and n = 48.



Fig. 1308: Histogram of release overhead under GSN-EDF, with m =48 and n = 48.



Fig. 1309: Histogram of release latency under GSN-EDF, with m = 48and n = 48.

min=0.15us max=17.22us avg=0.68us median=0.20us stdev=1.47us samples: total=3965931 fraction of samples 0.1 0.01 0.001 0.0001 0.50 2.50 4.50 6.50 8.50 10.50 12.50 14.50 16.50 overhead in microseconds (bin size = 1.00us)

Fig. 1310: Histogram of tick overhead under GSN-EDF, with m = 48and n = 48.













fraction of samples



Fig. 1314: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 48 and n = 48.



Fig. 1315: Histogram of release overhead under GSN-EDF-DI, with m = 48 and n = 48.



Fig. 1316: Histogram of release latency under GSN-EDF-DI, with m =48 and n = 48.



Fig. 1317: Histogram of tick overhead under GSN-EDF-DI, with m =48 and n = 48.

min=0.36us max=14.57us avg=1.96us median=1.91us stdev=0.63us samples: total=291179 fraction of samples 0. 0.01 0.00 0.0001 0.45 2.25 4.05 5.85 7.65 9.45 11.25 13.05 14.85 overhead in microseconds (bin size = 0.90us)

Fig. 1318: Histogram of context-switching overhead under GSN-EDF-DI, with m = 48 and n = 48.



Fig. 1319: Histogram of IPI latency under GSN-EDF-DI, with m = 48and n = 48.











Fig. 1322: Histogram of release overhead under G-EDF-MP, with m =48 and n = 48.



Fig. 1323: Histogram of release latency under G-EDF-MP, with m=48and n = 48.



Fig. 1324: Histogram of tick overhead under G-EDF-MP, with m = 48and n = 48.



Fig. 1325: Histogram of context-switching overhead under G-EDF-MP, with m = 48 and n = 48.

min=0.75us max=38.15us avg=3.75us median=2.83us stdev=2.62us samples: total=38833 fraction of samples 0.1 0.01 0.001 0.0001 1.50 7.50 13.50 19.50 25.50 31.50 37.50 overhead in microseconds (bin size = 3.00us)

Fig. 1326: Histogram of IPI latency under G-EDF-MP, with m=48and n = 48.



Fig. 1327: Histogram of client request latency under G-EDF-MP, with m = 48 and n = 48.









179



Fig. 1330: Histogram of scheduling overhead (b) under SD, with m = 48 and n = 48.



Fig. 1331: Histogram of release overhead under SD, with m = 48 and n = 48.



Fig. 1332: Histogram of release latency under SD, with m = 48 and n = 48.



Fig. 1333: Histogram of tick overhead under SD, with m = 48 and n = 48.

min=0.26us max=12.20us avg=1.37us median=1.28us stdev=0.64us samples: total=320833 fraction of samples 0. 0.01 0.00 0.0001 0.35 1.75 3.15 4.55 5.95 7.35 8.75 10.15 11.55 overhead in microseconds (bin size = 0.70us)











min=0.07us max=5.29us avg=0.09us median=0.07us stdev=0.10us samples: total=593323 fraction of samples 0.1 0.01 0.001 0.0001 0.15 0.75 1 35 1.95 2.55 3.15 3.75 4.35 4.95 overhead in microseconds (bin size = 0.30us)



n = 48.


Fig. 1338: Histogram of release overhead under GSN-EDF, with m =48 and n = 96.



Fig. 1339: Histogram of release latency under GSN-EDF, with m = 48and n = 96.



Fig. 1340: Histogram of tick overhead under GSN-EDF, with m = 48and n = 96.



Fig. 1341: Histogram of context-switching overhead under GSN-EDF, with m = 48 and n = 96.

samples: total=315609 fraction of samples 0.1 0.01 0.001 0.0001 1.50 7.50 13.50 19.50 25.50 31.50 37.50 43.50 49.50 overhead in microseconds (bin size = 3.00us)



min=0.22us max=381.26us avg=38.25us median=25.56us stdev=40.67us



Fig. 1343: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 48 and n = 96.









min=0.58us max=50.45us avg=2.11us median=1.73us stdev=1.31us



Fig. 1346: Histogram of release latency under GSN-EDF-DI, with m =48 and n = 96.



Fig. 1347: Histogram of tick overhead under GSN-EDF-DI, with m =48 and n = 96.



Fig. 1348: Histogram of context-switching overhead under GSN-EDF-DI, with m = 48 and n = 96.



Fig. 1349: Histogram of IPI latency under GSN-EDF-DI, with m = 48and n = 96.

min=0.20us max=41.24us avg=2.44us median=2.01us stdev=1.39us samples: total=594061 fraction of samples 0.1 0.01 0.001 0.0001 1.50 7.50 13.50 19.50 25.50 31.50 37.50 overhead in microseconds (bin size = 3.00us)











min=6.89us max=86.03us avg=32.31us median=31.21us stdev=10.64us samples: total=91558 fraction of samples 0.1 0.01 0.001 0.0001 2.50 12.50 22.50 32.50 42.50 52.50 62.50 72.50 82 50 overhead in microseconds (bin size = 5.00us)





Fig. 1354: Histogram of tick overhead under G-EDF-MP, with m = 48 and n = 96.



Fig. 1355: Histogram of context-switching overhead under G-EDF-MP, with m = 48 and n = 96.



Fig. 1356: Histogram of IPI latency under G-EDF-MP, with m = 48 and n = 96.



Fig. 1357: Histogram of client request latency under G-EDF-MP, with m=48 and n=96.

samples: total=595131 fraction of samples 0.1 0.01 0.001 0.0001 2.50 12.50 22.50 32.50 42.50 52.50 62.50 72.50 82.50 overhead in microseconds (bin size = 5.00us)

Fig. 1358: Histogram of DSP handler overhead under G-EDF-MP, with m = 48 and n = 96.











Fig. 1361: Histogram of release overhead under SD, with m = 48 and n = 96.

min=0.07us max=82.65us avg=3.52us median=3.31us stdev=2.18us



Fig. 1362: Histogram of release latency under SD, with m = 48 and n = 96.



Fig. 1363: Histogram of tick overhead under SD, with m = 48 and n = 96.



Fig. 1364: Histogram of context-switching overhead under SD, with m = 48 and n = 96.



Fig. 1365: Histogram of IPI latency under SD, with m = 48 and n = 96.

min=0.33us max=482.83us avg=68.36us median=45.91us stdev=70.25us



Fig. 1366: Histogram of scheduling overhead (a) under GSN-EDF, with m=48 and n=144.









min=5.36us max=86.33us avg=33.94us median=32.48us stdev=11.10us samples: total=94570 fraction of samples 0.1 0.01 0.00 0.0001 2.50 12.50 22.50 32.50 42.50 52.50 62.50 72.50 82.50 overhead in microseconds (bin size = 5.00us)





min=0.15us max=21.81us avg=0.82us median=0.26us stdev=1.61us

Fig. 1370: Histogram of tick overhead under GSN-EDF, with m = 48and n = 144.



Fig. 1371: Histogram of context-switching overhead under GSN-EDF, with m = 48 and n = 144.



Fig. 1372: Histogram of IPI latency under GSN-EDF, with m = 48 and n = 144.



Fig. 1373: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 48 and n = 144.

min=0.07us max=9.01us avg=0.10us median=0.07us stdev=0.12us samples: total=905737 fraction of samples 0.1 0.01 0.001 0.0001 0.30 1.50 2.70 3.90 5.10 6.30 7.50 8.70

Fig. 1374: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 48 and n = 144.



overhead in microseconds (bin size = 20.00us) Fig. 1375: Histogram of release overhead under GSN-EDF-DI, with m = 48 and n = 144.









overhead in microseconds (bin size = 0.60us)



Fig. 1378: Histogram of context-switching overhead under GSN-EDF-DI, with m = 48 and n = 144.



Fig. 1379: Histogram of IPI latency under GSN-EDF-DI, with m = 48and n = 144.



Fig. 1380: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 48 and n = 144.



Fig. 1381: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 48 and n = 144.

min=0.73us max=159.86us avg=35.11us median=30.21us stdev=23.57us samples: total=94564 fraction of samples 0. 0.01 0.001 0.0001 5.00 25.00 45.00 65.00 85.00 105.00 125.00 145.00

overhead in microseconds (bin size = 10.00us) Fig. 1382: Histogram of release overhead under G-EDF-MP, with m =

min=6.90us max=87.03us avg=33.65us median=31.89us stdev=11.14us



Fig. 1383: Histogram of release latency under G-EDF-MP, with m = 48and n = 144.





min=0.29us max=18.25us avg=2.00us median=1.89us stdev=0.86us samples: total=902856 fraction of samples 0.1 0.01 0.001 0.0001 1.00 3.00 5.00 7.00 9.00 11.00 13.00 15.00 17.00 19.00 overhead in microseconds (bin size = 2.00us)



48 and n = 144.

min=0.20us max=20.39us avg=2.17us median=1.82us stdev=1.21us



Fig. 1386: Histogram of IPI latency under G-EDF-MP, with m = 48 and n = 144.



Fig. 1387: Histogram of client request latency under G-EDF-MP, with m = 48 and n = 144.



Fig. 1388: Histogram of DSP handler overhead under G-EDF-MP, with m = 48 and n = 144.



Fig. 1389: Histogram of scheduling overhead (a) under SD, with m = 48 and n = 144.

min=0.06us max=634.65us avg=4.77us median=0.08us stdev=11.52us 0.01 0.001 0.001 0.001 0.001 0.001 0.000 100.00 180.00 260.00 340.00 420.00 500.00 580.00 overhead in microseconds (bin size = 40.00us)

Fig. 1390: Histogram of scheduling overhead (b) under SD, with m = 48 and n = 144.



overhead in microseconds (bin size = 20.00us) Fig. 1391: Histogram of release overhead under SD, with m = 48 and n = 144.

170.00

210.00

250.00

290.00

130.00





min=0.16us max=77.66us avg=1.49us median=0.25us stdev=2.84us samples: total=3232464 fraction of samples 0.1 0.01 0.001 0.0001 2.50 12.50 22.50 32 50 42.50 52.50 62.50 72.50 overhead in microseconds (bin size = 5.00us)



0.0001

10.00

50.00

90.00

min=0.22us max=313.30us avg=22.31us median=12.24us stdev=27.01us



Fig. 1394: Histogram of context-switching overhead under SD, with m = 48 and n = 144.



Fig. 1395: Histogram of IPI latency under SD, with m = 48 and n = 144.



Fig. 1396: Histogram of scheduling overhead (a) under GSN-EDF, with m = 48 and n = 192.



Fig. 1397: Histogram of scheduling overhead (b) under GSN-EDF, with m = 48 and n = 192.

min=1.08us max=463.10us avg=91.00us median=62.98us stdev=76.63us samples: total=95628 fraction of samples 0. 0.01 0.001 0.0001 15.00 75.00 135.00 195.00 255.00 315.00 375.00 435.00

overhead in microseconds (bin size = 30.00us) Fig. 1398: Histogram of release overhead under GSN-EDF, with m =48 and n = 192.



overhead in microseconds (bin size = 5.00us) Fig. 1399: Histogram of release latency under GSN-EDF, with m=48

42.50 52.50

32.50





min=0.37us max=14.08us avg=2.07us median=2.04us stdev=0.54us samples, total=1166127 fraction of samples 0. 0.01 0.001 0.0001 0.40 2.00 3.60 5.20 6.80 8.40 10.00 11.60 13.20 overhead in microseconds (bin size = 0.80us)



and n = 192.

2.50

12.50

22.50

fraction of samples

62.50

72.50 82.50



Fig. 1402: Histogram of IPI latency under GSN-EDF, with m=48 and n=192.



Fig. 1403: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 48 and n = 192.



Fig. 1404: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 48 and n = 192.



Fig. 1405: Histogram of release overhead under GSN-EDF-DI, with m = 48 and n = 192.

min=7.09us max=74.36us avg=33.92us median=32.19us stdev=9.59us samples: total=95633 fraction of samples 0.1 0.01 0.001 0.0001 2.50 12.50 22.50 32.50 42.50 52.50 62.50 72.50 overhead in microseconds (bin size = 5.00us)

Fig. 1406: Histogram of release latency under GSN-EDF-DI, with m = 48 and n = 192.







Fig. 1408: Histogram of context-switching overhead under GSN-EDF-DI, with m=48 and n=192.



Fig. 1409: Histogram of IPI latency under GSN-EDF-DI, with m=48 and n=192.



Fig. 1410: Histogram of scheduling overhead (a) under G-EDF-MP, with m=48 and n=192.



Fig. 1411: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 48 and n = 192.



Fig. 1412: Histogram of release overhead under G-EDF-MP, with m = 48 and n = 192.



Fig. 1413: Histogram of release latency under G-EDF-MP, with m=48 and n=192.

min=0.15us max=15.03us avg=0.57us median=0.19us stdev=1.28us



Fig. 1414: Histogram of tick overhead under G-EDF-MP, with m = 48 and n = 192.



Fig. 1415: Histogram of context-switching overhead under G-EDF-MP, with m = 48 and n = 192.







Fig. 1417: Histogram of client request latency under G-EDF-MP, with m=48 and n=192.



Fig. 1418: Histogram of DSP handler overhead under G-EDF-MP, with m = 48 and n = 192.



Fig. 1419: Histogram of scheduling overhead (a) under SD, with m = 48 and n = 192.



Fig. 1420: Histogram of scheduling overhead (b) under SD, with m = 48 and n = 192.



Fig. 1421: Histogram of release overhead under SD, with m=48 and n=192.

min=1.42us max=488.16us avg=17.73us median=8.55us stdev=33.05us



Fig. 1422: Histogram of release latency under SD, with m=48 and n=192.











Fig. 1425: Histogram of IPI latency under SD, with m = 48 and n = 192.



Fig. 1426: Histogram of scheduling overhead (a) under GSN-EDF, with m = 48 and n = 240.



Fig. 1427: Histogram of scheduling overhead (b) under GSN-EDF, with m = 48 and n = 240.



Fig. 1428: Histogram of release overhead under GSN-EDF, with m =48 and n = 240.



Fig. 1429: Histogram of release latency under GSN-EDF, with m = 48and n = 240.

min=0.15us max=17.26us avg=0.94us median=0.34us stdev=1.72us samples: total=3578766 fraction of samples 0.1 0.01 0.001 0.0001 0.50 2.50 4.50 6.50 8.50 12.50 14.50 16.50 overhead in microseconds (bin size = 1.00us)

Fig. 1430: Histogram of tick overhead under GSN-EDF, with m = 48and n = 240.















Fig. 1434: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 48 and n = 240.



Fig. 1435: Histogram of release overhead under GSN-EDF-DI, with m = 48 and n = 240.



Fig. 1436: Histogram of release latency under GSN-EDF-DI, with m =48 and n = 240.



Fig. 1437: Histogram of tick overhead under GSN-EDF-DI, with m =48 and n = 240.

min=0.36us max=15.84us avg=1.99us median=1.96us stdev=0.56us samples: total=1514305 fraction of samples 0.1 0.01 0.00 0.0001 0.45 2.25 4.05 5.85 7.65 9.45 11.25 13.05 14.85 overhead in microseconds (bin size = 0.90us)

Fig. 1438: Histogram of context-switching overhead under GSN-EDF-DI, with m = 48 and n = 240.



Fig. 1439: Histogram of IPI latency under GSN-EDF-DI, with m = 48and n = 240.











Fig. 1442: Histogram of release overhead under G-EDF-MP, with m = 48 and n = 240.



Fig. 1443: Histogram of release latency under G-EDF-MP, with m=48 and n=240.



Fig. 1444: Histogram of tick overhead under G-EDF-MP, with m = 48 and n = 240.



Fig. 1445: Histogram of context-switching overhead under G-EDF-MP, with m=48 and n=240.

min=0.58us max=43.42us avg=2.11us median=1.78us stdev=1.30us samples: total=880152 fraction of samples 0.1 0.01 0.001 0.0001 1.50 7.50 13.50 19.50 25.50 31.50 37.50 43.50 overhead in microseconds (bin size = 3.00us)

Fig. 1446: Histogram of IPI latency under G-EDF-MP, with m = 48 and n = 240.



Fig. 1447: Histogram of client request latency under G-EDF-MP, with m = 48 and n = 240.







Fig. 1449: Histogram of scheduling overhead (a) under SD, with m = 48 and n = 240.



Fig. 1450: Histogram of scheduling overhead (b) under SD, with m = 48 and n = 240.



Fig. 1451: Histogram of release overhead under SD, with m = 48 and n = 240.



Fig. 1452: Histogram of release latency under SD, with m = 48 and n = 240.



Fig. 1453: Histogram of tick overhead under SD, with m = 48 and n = 240.

min=0.23us max=21.22us avg=1.77us median=1.72us stdev=0.85us



Fig. 1454: Histogram of context-switching overhead under SD, with m=48 and n=240.







Fig. 1456: Histogram of scheduling overhead (a) under GSN-EDF, with m=48 and n=288.







Fig. 1458: Histogram of release overhead under GSN-EDF, with m = 48 and n = 288.



Fig. 1459: Histogram of release latency under GSN-EDF, with m = 48 and n = 288.



Fig. 1460: Histogram of tick overhead under GSN-EDF, with m = 48 and n = 288.



Fig. 1461: Histogram of context-switching overhead under GSN-EDF, with m = 48 and n = 288.



Fig. 1462: Histogram of IPI latency under GSN-EDF, with m = 48 and n = 288.









Fig. 1464: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 48 and n = 288.



Fig. 1465: Histogram of release overhead under GSN-EDF-DI, with m=48 and n=288.



Fig. 1466: Histogram of release latency under GSN-EDF-DI, with m = 48 and n = 288.



Fig. 1467: Histogram of tick overhead under GSN-EDF-DI, with m = 48 and n = 288.



Fig. 1468: Histogram of context-switching overhead under GSN-EDF-DI, with m = 48 and n = 288.



Fig. 1469: Histogram of IPI latency under GSN-EDF-DI, with m = 48 and n = 288.

min=0.18us max=49.14us avg=2.05us median=1.78us stdev=1.07us samples: total=1870654 fraction of samples 0.1 0.01 0.001 0.0001 1.50 7.50 13.50 19.50 25.50 31.50 37.50 43.50 49.50 overhead in microseconds (bin size = 3.00us)









Fig. 1472: Histogram of release overhead under G-EDF-MP, with m = 48 and n = 288.



Fig. 1473: Histogram of release latency under G-EDF-MP, with m = 48 and n = 288.



Fig. 1474: Histogram of tick overhead under G-EDF-MP, with m = 48and n = 288.



Fig. 1475: Histogram of context-switching overhead under G-EDF-MP, with m = 48 and n = 288.



Fig. 1476: Histogram of IPI latency under G-EDF-MP, with m = 48and n = 288.



Fig. 1477: Histogram of client request latency under G-EDF-MP, with m = 48 and n = 288.

min=0.07us max=86.15us avg=3.85us median=3.28us stdev=2.29us samples: total=1747507 fraction of samples 0.1 0.01 0.001 0.0001 2.50 12.50 22.50 32.50 42.50 52.50 62.50 72.50 82.50 overhead in microseconds (bin size = 5.00us)

Fig. 1478: Histogram of DSP handler overhead under G-EDF-MP, with m = 48 and n = 288.















Fig. 1482: Histogram of release latency under SD, with m = 48 and n = 288.



Fig. 1483: Histogram of tick overhead under SD, with m = 48 and n = 288.



Fig. 1484: Histogram of context-switching overhead under SD, with m = 48 and n = 288.



Fig. 1485: Histogram of IPI latency under SD, with m=48 and n=288.

min=0.20us max=796.57us avg=136.20us median=113.79us stdev=108.79us





min=0.07us max=7.60us avg=0.10us median=0.07us stdev=0.15us amples: total=2252958 0.01 0.001 0.001 0.001 0.25 1.25 2.25 3.25 4.25 5.25 6.25 7.25 overhead in microseconds (bin size = 0.50us)





Fig. 1488: Histogram of release overhead under GSN-EDF, with m = 48 and n = 336.



Fig. 1489: Histogram of release latency under GSN-EDF, with m=48 and n=336.



Fig. 1490: Histogram of tick overhead under GSN-EDF, with m = 48 and n = 336.



Fig. 1491: Histogram of context-switching overhead under GSN-EDF, with m = 48 and n = 336.



Fig. 1492: Histogram of IPI latency under GSN-EDF, with m = 48 and n = 336.



Fig. 1493: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m=48 and n=336.

samples: total=2264374 fraction of samples 0.1 0.01 0.00 0.0001 0.20 1.00 1.80 2.60 3.40 4.20 5.00 5.80 6.60 overhead in microseconds (bin size = 0.40us)



min=1.16us max=495.71us avg=128.38us median=105.86us stdev=87.86us



Fig. 1495: Histogram of release overhead under GSN-EDF-DI, with m=48 and n=336.





min=0.15us max=17.42us avg=1.05us median=0.41us stdev=1.82us samples, total=3492428 fraction of samples 0.1 0.01 0.001 0.0001 0.50 2.50 4.50 6.50 8.50 10.50 12.50 14.50 16.50 overhead in microseconds (bin size = 1.00us)



min=0.07us max=7.14us avg=0.10us median=0.07us stdev=0.15us



Fig. 1498: Histogram of context-switching overhead under GSN-EDF-DI, with m = 48 and n = 336.



Fig. 1499: Histogram of IPI latency under GSN-EDF-DI, with m = 48 and n = 336.



Fig. 1500: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 48 and n = 336.



Fig. 1501: Histogram of scheduling overhead (b) under G-EDF-MP, with m=48 and n=336.

min=0.78us max=190.89us avg=61.83us median=56.52us stdev=36.01us



overhead in microseconds (bin size = 15.00us) Fig. 1502: Histogram of release overhead under G-EDF-MP, with m = 48 and n = 336.

min=6.92us max=87.32us avg=31.51us median=29.78us stdev=12.67us



Fig. 1503: Histogram of release latency under G-EDF-MP, with m = 48 and n = 336.





min=0.28us max=20.54us avg=2.01us median=1.95us stdev=0.77us amples: total=2248156 0.01 0.001 0.001 1.00 5.00 9.00 13.00 17.00 21.00 overhead in microseconds (bin size = 2.00us)





Fig. 1506: Histogram of IPI latency under G-EDF-MP, with m=48 and n=336.



Fig. 1507: Histogram of client request latency under G-EDF-MP, with m = 48 and n = 336.



Fig. 1508: Histogram of DSP handler overhead under G-EDF-MP, with m = 48 and n = 336.



Fig. 1509: Histogram of scheduling overhead (a) under SD, with m = 48 and n = 336.

min=0.06us max=791.65us avg=9.01us median=0.08us stdev=23.19us amples: total=2401596 0.01 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.0000 0.000 0.000 0.000 0.000 0.000 0.000

overhead in microseconds (bin size = 50.00us) Fig. 1510: Histogram of scheduling overhead (b) under SD, with m = 48 and n = 336.













Fig. 1513: Histogram of tick overhead under SD, with m = 48 and n = 336.



Fig. 1514: Histogram of context-switching overhead under SD, with m = 48 and n = 336.



Fig. 1515: Histogram of IPI latency under SD, with m = 48 and n = 336.



Fig. 1516: Histogram of scheduling overhead (a) under GSN-EDF, with m = 48 and n = 384.



Fig. 1517: Histogram of scheduling overhead (b) under GSN-EDF, with m=48 and n=384.

min=1.06us max=560.32us avg=164.97us median=133.69us stdev=114.13us

20.00 100.00 180.00 260.00 340.00 420.00 500.00 580.00 overhead in microseconds (bin size = 40.00us)

Fig. 1518: Histogram of release overhead under GSN-EDF, with m = 48 and n = 384.



Fig. 1519: Histogram of release latency under GSN-EDF, with m = 48 and n = 384.



Fig. 1520: Histogram of tick overhead under GSN-EDF, with m = 48 and n = 384.



Fig. 1521: Histogram of context-switching overhead under GSN-EDF, with m=48 and n=384.



Fig. 1522: Histogram of IPI latency under GSN-EDF, with m = 48 and n = 384.



Fig. 1523: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 48 and n = 384.



Fig. 1524: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 48 and n = 384.



Fig. 1525: Histogram of release overhead under GSN-EDF-DI, with m = 48 and n = 384.

min=2.86us max=76.99us avg=32.55us median=31.73us stdev=9.72us samples: total=97914 fraction of samples 0.1 0.01 0.001 0.0001 2.50 12.50 22.50 32.50 42.50 52.50 62.50 72.50 overhead in microseconds (bin size = 5.00us)

Fig. 1526: Histogram of release latency under GSN-EDF-DI, with m =48 and n = 384.







Fig. 1528: Histogram of context-switching overhead under GSN-EDF-DI, with m = 48 and n = 384.





48 and n = 384.



Fig. 1530: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 48 and n = 384.



Fig. 1531: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 48 and n = 384.



Fig. 1532: Histogram of release overhead under G-EDF-MP, with m =48 and n = 384.



Fig. 1533: Histogram of release latency under G-EDF-MP, with m = 48and n = 384.

min=0.15us max=15.04us avg=0.78us median=0.26us stdev=1.55us samples: total=2517403



Fig. 1534: Histogram of tick overhead under G-EDF-MP, with m = 48and n = 384.



Fig. 1535: Histogram of context-switching overhead under G-EDF-MP, with m = 48 and n = 384.











Fig. 1538: Histogram of DSP handler overhead under G-EDF-MP, with m = 48 and n = 384.



Fig. 1539: Histogram of scheduling overhead (a) under SD, with m = 48 and n = 384.



Fig. 1540: Histogram of scheduling overhead (b) under SD, with m = 48 and n = 384.



Fig. 1541: Histogram of release overhead under SD, with m=48 and n=384.

min=1.50us max=946.98us avg=37.48us median=12.83us stdev=71.50us



Fig. 1542: Histogram of release latency under SD, with m = 48 and n = 384.



Fig. 1543: Histogram of tick overhead under SD, with m = 48 and n = 384.







Fig. 1545: Histogram of IPI latency under SD, with m = 48 and n = 384.



Fig. 1546: Histogram of scheduling overhead (a) under GSN-EDF, with m = 48 and n = 432.



Fig. 1547: Histogram of scheduling overhead (b) under GSN-EDF, with m = 48 and n = 432.



Fig. 1548: Histogram of release overhead under GSN-EDF, with m =48 and n = 432.



Fig. 1549: Histogram of release latency under GSN-EDF, with m = 48and n = 432.

samples: total=3727377 fraction of samples 0.1 0.01 0.00 0.0001 1.00 5.00 9.00 13.00 17.00 21.00 25.00 29.00 33.00 overhead in microseconds (bin size = 2.00us)

Fig. 1550: Histogram of tick overhead under GSN-EDF, with m = 48and n = 432.



Fig. 1551: Histogram of context-switching overhead under GSN-EDF, with m = 48 and n = 432.









min=0.15us max=33.28us avg=1.12us median=0.44us stdev=1.86us



Fig. 1554: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 48 and n = 432.



Fig. 1555: Histogram of release overhead under GSN-EDF-DI, with m = 48 and n = 432.



Fig. 1556: Histogram of release latency under GSN-EDF-DI, with m =48 and n = 432.



Fig. 1557: Histogram of tick overhead under GSN-EDF-DI, with m =48 and n = 432.

min=0.35us max=16.41us avg=1.96us median=1.91us stdev=0.54us samples: total=2902148 fraction of samples 0. 0.01 0.001 0.0001 0.50 2.50 4.50 6.50 8.50 10.50 12.50 14.50 16.50

overhead in microseconds (bin size = 1.00us) Fig. 1558: Histogram of context-switching overhead under GSN-EDF-DI, with m = 48 and n = 432.















Fig. 1562: Histogram of release overhead under G-EDF-MP, with m =48 and n = 432.



Fig. 1563: Histogram of release latency under G-EDF-MP, with m=48and n = 432.



Fig. 1564: Histogram of tick overhead under G-EDF-MP, with m = 48and n = 432.



Fig. 1565: Histogram of context-switching overhead under G-EDF-MP, with m = 48 and n = 432.

min=0.62us max=47.97us avg=1.89us median=1.65us stdev=1.06us samples: total=2173864 fraction of samples 0.1 0.01 0.001 0.0001 1.50 7.50 13.50 19.50 25.50 31.50 37.50 43.50 overhead in microseconds (bin size = 3.00us)

Fig. 1566: Histogram of IPI latency under G-EDF-MP, with m = 48and n = 432.



Fig. 1567: Histogram of client request latency under G-EDF-MP, with m = 48 and n = 432.











Fig. 1570: Histogram of scheduling overhead (b) under SD, with m = 48 and n = 432.



Fig. 1571: Histogram of release overhead under SD, with m = 48 and n = 432.



Fig. 1572: Histogram of release latency under SD, with m = 48 and n = 432.



Fig. 1573: Histogram of tick overhead under SD, with m=48 and n=432.

min=0.21us max=22.45us avg=1.82us median=1.75us stdev=0.94us



Fig. 1574: Histogram of context-switching overhead under SD, with m = 48 and n = 432.



Fig. 1575: Histogram of IPI latency under SD, with m = 48 and n = 432.



Fig. 1576: Histogram of scheduling overhead (a) under GSN-EDF, with m = 48 and n = 480.



Fig. 1577: Histogram of scheduling overhead (b) under GSN-EDF, with m=48 and n=480.



overhead in microseconds (bin size = 40.00us) Fig. 1578: Histogram of release overhead under GSN-EDF, with m = 48 and n = 480.



Fig. 1579: Histogram of release latency under GSN-EDF, with m = 48 and n = 480.



Fig. 1580: Histogram of tick overhead under GSN-EDF, with m = 48 and n = 480.



Fig. 1581: Histogram of context-switching overhead under GSN-EDF, with m=48 and n=480.

min=0.53us max=425.43us avg=1.96us median=1.64us stdev=1.62us 0.01 0.001 0.001 12.50 62.50 112.50 162.50 212.50 262.50 312.50 362.50 412.50 overhead in microseconds (bin size = 25.00us)



min=0.20us max=942.79us avg=134.68us median=120.19us stdev=108.92us







Fig. 1584: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 48 and n = 480.



Fig. 1585: Histogram of release overhead under GSN-EDF-DI, with m=48 and n=480.



Fig. 1586: Histogram of release latency under GSN-EDF-DI, with m = 48 and n = 480.



Fig. 1587: Histogram of tick overhead under GSN-EDF-DI, with m = 48 and n = 480.



Fig. 1588: Histogram of context-switching overhead under GSN-EDF-DI, with m = 48 and n = 480.



Fig. 1589: Histogram of IPI latency under GSN-EDF-DI, with m = 48 and n = 480.

Fig. 1590: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 48 and n = 480.







Fig. 1592: Histogram of release overhead under G-EDF-MP, with m = 48 and n = 480.



Fig. 1593: Histogram of release latency under G-EDF-MP, with m=48 and n=480.

min=0.17us max=25.36us avg=1.83us median=1.67us stdev=0.84us



Fig. 1594: Histogram of tick overhead under G-EDF-MP, with m = 48 and n = 480.



Fig. 1595: Histogram of context-switching overhead under G-EDF-MP, with m = 48 and n = 480.



Fig. 1596: Histogram of IPI latency under G-EDF-MP, with m = 48 and n = 480.



Fig. 1597: Histogram of client request latency under G-EDF-MP, with m=48 and n=480.

min=0.07us max=85.69us avg=4.83us median=3.62us stdev=3.04us



Fig. 1598: Histogram of DSP handler overhead under G-EDF-MP, with m = 48 and n = 480.









min=0.19us max=908.21us avg=50.73us median=26.89us stdev=63.75us amples: total=2650962 0.01 0.001 0.001 0.001 50.00 150.00 250.00 350.00 450.00 550.00 650.00 750.00 850.00 950.00 overhead in microseconds (bin size = 100.00us)





Fig. 1602: Histogram of release latency under SD, with m = 48 and n = 480.



Fig. 1603: Histogram of tick overhead under SD, with m = 48 and n = 480.



Fig. 1604: Histogram of context-switching overhead under SD, with m = 48 and n = 480.



Fig. 1605: Histogram of IPI latency under SD, with m = 48 and n = 480.

min=0.25us max=856.64us avg=31.87us median=11.93us stdev=47.58us



Fig. 1606: Histogram of scheduling overhead (a) under GSN-EDF, with m=64 and n=64.



Fig. 1607: Histogram of scheduling overhead (b) under GSN-EDF, with m = 64 and n = 64.



Fig. 1608: Histogram of release overhead under GSN-EDF, with m = 64 and n = 64.



Fig. 1609: Histogram of release latency under GSN-EDF, with m=64 and n=64.



Fig. 1610: Histogram of tick overhead under GSN-EDF, with m = 64and n = 64.



Fig. 1611: Histogram of context-switching overhead under GSN-EDF, with m = 64 and n = 64.



Fig. 1612: Histogram of IPI latency under GSN-EDF, with m = 64 and n = 64.



Fig. 1613: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 64 and n = 64.



Fig. 1614: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 64 and n = 64.



Fig. 1615: Histogram of release overhead under GSN-EDF-DI, with m = 64 and n = 64.











Fig. 1618: Histogram of context-switching overhead under GSN-EDF-DI, with m = 64 and n = 64.



Fig. 1619: Histogram of IPI latency under GSN-EDF-DI, with m = 64 and n = 64.



Fig. 1620: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 64 and n = 64.



Fig. 1621: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 64 and n = 64.

min=0.56us max=121.25us avg=26.91us median=22.50us stdev=18.61us



Fig. 1622: Histogram of release overhead under G-EDF-MP, with m = 64 and n = 64.

min=9.01us max=111.78us avg=38.73us median=36.15us stdev=13.20us



Fig. 1623: Histogram of release latency under G-EDF-MP, with m = 64 and n = 64.



Fig. 1624: Histogram of tick overhead under G-EDF-MP, with m = 64 and n = 64.



Fig. 1625: Histogram of context-switching overhead under G-EDF-MP, with m=64 and n=64.


Fig. 1626: Histogram of IPI latency under G-EDF-MP, with m=64 and n=64.



Fig. 1627: Histogram of client request latency under G-EDF-MP, with m = 64 and n = 64.



Fig. 1628: Histogram of DSP handler overhead under G-EDF-MP, with m = 64 and n = 64.



Fig. 1629: Histogram of scheduling overhead (a) under SD, with m = 64 and n = 64.

min=0.06us max=246.23us avg=0.77us median=0.07us stdev=3.75us amples: tdtal=423396 0.01 0.001 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.000 0.0000 0.0000 0.0000 0.0000 0.0000 0.0000 0.0000 0.0000 0.0000 0.0000 0.0000 0.0000 0.0000 0.0000 0.0000 0.0000 0.0000 0.0000 0.0000 0.0000 0.0000 0.0000 0.00000 0.0000 0.0000 0.0000 0.0000 0

Fig. 1630: Histogram of scheduling overhead (b) under SD, with m = 64 and n = 64.





Fig. 1631: Histogram of release overhead under SD, with m = 64 and n = 64.







Fig. 1633: Histogram of tick overhead under SD, with m = 64 and n = 64.



Fig. 1634: Histogram of context-switching overhead under SD, with m = 64 and n = 64.



Fig. 1635: Histogram of IPI latency under SD, with m = 64 and n = 64.



Fig. 1636: Histogram of scheduling overhead (a) under GSN-EDF, with m = 64 and n = 128.



Fig. 1637: Histogram of scheduling overhead (b) under GSN-EDF, with m=64 and n=128.

min=0.51us max=693.65us avg=86.25us median=62.87us stdev=72.97us samples: total=92920 0.01

0.001 0.0001 20.00 100.00 180.00 260.00 340.00 420.00 500.00 580.00 660.00

overhead in microseconds (bin size = 40.00us) Fig. 1638: Histogram of release overhead under GSN-EDF, with m = 64 and n = 128.

min=6.41us max=114.51us avg=38.50us median=36.48us stdev=13.42us



Fig. 1639: Histogram of release latency under GSN-EDF, with m = 64 and n = 128.



Fig. 1640: Histogram of tick overhead under GSN-EDF, with m = 64 and n = 128.



Fig. 1641: Histogram of context-switching overhead under GSN-EDF, with m = 64 and n = 128.



Fig. 1642: Histogram of IPI latency under GSN-EDF, with m = 64 and n = 128.



Fig. 1643: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 64 and n = 128.



Fig. 1644: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 64 and n = 128.



Fig. 1645: Histogram of release overhead under GSN-EDF-DI, with m = 64 and n = 128.

min=9.06us max=104.55us avg=45.35us median=42.52us stdev=13.64us 0.1 0.01 0.001 0.001



65.00

85.00

105.00

45.00



overhead in microseconds (bin size = 2.00us) Fig. 1647: Histogram of tick overhead under GSN-EDF-DI, with m = 64 and n = 128.



Fig. 1648: Histogram of context-switching overhead under GSN-EDF-DI, with m = 64 and n = 128.





ction

fraction of samples

0.0001

5.00

25.00



Fig. 1650: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 64 and n = 128.



Fig. 1651: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 64 and n = 128.



Fig. 1652: Histogram of release overhead under G-EDF-MP, with m = 64 and n = 128.



Fig. 1653: Histogram of release latency under G-EDF-MP, with m=64 and n=128.

min=0.15us max=17.15us avg=0.88us median=0.24us stdev=2.05us samples: total=5347796 fraction of samples 0.1 0.01 0.001 0.0001 0.50 2.50 4.50 6.50 8.50 10.50 12.50 14.50 16.50

overhead in microseconds (bin size = 1.00us) Fig. 1654: Histogram of tick overhead under G-EDF-MP, with m = 64 and n = 128.















Fig. 1658: Histogram of DSP handler overhead under G-EDF-MP, with m = 64 and n = 128.



Fig. 1659: Histogram of scheduling overhead (a) under SD, with m = 64 and n = 128.



Fig. 1660: Histogram of scheduling overhead (b) under SD, with m = 64 and n = 128.



Fig. 1661: Histogram of release overhead under SD, with m = 64 and n = 128.

min=1.46us max=392.58us avg=10.90us median=6.62us stdev=18.75us



Fig. 1662: Histogram of release latency under SD, with m = 64 and n = 128.











Fig. 1665: Histogram of IPI latency under SD, with m = 64 and n = 128.



Fig. 1666: Histogram of scheduling overhead (a) under GSN-EDF, with m = 64 and n = 192.



Fig. 1667: Histogram of scheduling overhead (b) under GSN-EDF, with m=64 and n=192.



Fig. 1668: Histogram of release overhead under GSN-EDF, with m = 64 and n = 192.



Fig. 1669: Histogram of release latency under GSN-EDF, with m=64 and n=192.



Fig. 1670: Histogram of tick overhead under GSN-EDF, with m = 64 and n = 192.



Fig. 1671: Histogram of context-switching overhead under GSN-EDF, with m = 64 and n = 192.









samples: total=97083



Fig. 1674: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 64 and n = 192.



Fig. 1675: Histogram of release overhead under GSN-EDF-DI, with m=64 and n=192.



Fig. 1676: Histogram of release latency under GSN-EDF-DI, with m = 64 and n = 192.



Fig. 1677: Histogram of tick overhead under GSN-EDF-DI, with m = 64 and n = 192.

min=0.36us max=20.28us avg=2.20us median=2.10us stdev=0.80us



Fig. 1678: Histogram of context-switching overhead under GSN-EDF-DI, with m = 64 and n = 192.







Fig. 1680: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 64 and n = 192.







Fig. 1682: Histogram of release overhead under G-EDF-MP, with m = 64 and n = 192.



Fig. 1683: Histogram of release latency under G-EDF-MP, with m=64 and n=192.



Fig. 1684: Histogram of tick overhead under G-EDF-MP, with m = 64 and n = 192.



Fig. 1685: Histogram of context-switching overhead under G-EDF-MP, with m=64 and n=192.

min=0.58us max=58.32us avg=2.66us median=2.03us stdev=2.06us samples: total=445467 fraction of samples 0.1 0.01 0.001 0.0001 2.00 10.00 18.00 26.00 34.00 42.00 50.00 58.00 overhead in microseconds (bin size = 4.00us)

Fig. 1686: Histogram of IPI latency under G-EDF-MP, with m = 64 and n = 192.



Fig. 1687: Histogram of client request latency under G-EDF-MP, with m = 64 and n = 192.



Fig. 1688: Histogram of DSP handler overhead under G-EDF-MP, with m = 64 and n = 192.



Fig. 1689: Histogram of scheduling overhead (a) under SD, with m = 64 and n = 192.



Fig. 1690: Histogram of scheduling overhead (b) under SD, with m = 64 and n = 192.



Fig. 1691: Histogram of release overhead under SD, with m = 64 and n = 192.



Fig. 1692: Histogram of release latency under SD, with m = 64 and n = 192.



Fig. 1693: Histogram of tick overhead under SD, with m=64 and n=192.

min=0.23us max=28.13us avg=1.82us median=1.71us stdev=1.14us samples: total=1378790 0.01 0.001 0.001 1.00 5.00 9.00 13.00 17.00 21.00 25.00 29.00

Fig. 1694: Histogram of context-switching overhead under SD, with m = 64 and n = 192.

overhead in microseconds (bin size = 2.00us)







Fig. 1696: Histogram of scheduling overhead (a) under GSN-EDF, with m=64 and n=256.







Fig. 1698: Histogram of release overhead under GSN-EDF, with m = 64 and n = 256.



Fig. 1699: Histogram of release latency under GSN-EDF, with m = 64 and n = 256.



Fig. 1700: Histogram of tick overhead under GSN-EDF, with m = 64 and n = 256.



Fig. 1701: Histogram of context-switching overhead under GSN-EDF, with m=64 and n=256.

min=0.47us max=419.33us avg=2.26us median=1.63us stdev=1.87us 0.01 0.001 0.001 12.50 62.50 112.50 162.50 212.50 262.50 312.50 362.50 412.50 overhead in microseconds (bin size = 25.00us)



min=0.20us max=791.85us avg=164.68us median=129.98us stdev=139.00us



Fig. 1703: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m=64 and n=256.







Fig. 1705: Histogram of release overhead under GSN-EDF-DI, with $m=64~{\rm and}~n=256.$



Fig. 1706: Histogram of release latency under GSN-EDF-DI, with m =64 and n = 256.



Fig. 1707: Histogram of tick overhead under GSN-EDF-DI, with m =64 and n = 256.



Fig. 1708: Histogram of context-switching overhead under GSN-EDF-DI, with m = 64 and n = 256.



Fig. 1709: Histogram of IPI latency under GSN-EDF-DI, with m = 64and n = 256.

min=0.17us max=26.97us avg=2.41us median=1.94us stdev=1.62us samples: total=1604770 fraction of samples 0.1 0.01 0.001 0.0001 1.00 5.00 9.00 13.00 17.00 21.00 25.00 overhead in microseconds (bin size = 2.00us)

Fig. 1710: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 64 and n = 256.







Fig. 1712: Histogram of release overhead under G-EDF-MP, with m =64 and n = 256.

min=9.32us max=114.19us avg=37.52us median=34.02us stdev=14.40us samples: total=97348 fraction of samples 0.1 0.01 0.001 0.0001 5.00 25.00 45.00 65.00 85.00 105.00 overhead in microseconds (bin size = 10.00us)





Fig. 1714: Histogram of tick overhead under G-EDF-MP, with m = 64and n = 256.



Fig. 1715: Histogram of context-switching overhead under G-EDF-MP, with m = 64 and n = 256.



Fig. 1716: Histogram of IPI latency under G-EDF-MP, with m = 64and n = 256.



Fig. 1717: Histogram of client request latency under G-EDF-MP, with m = 64 and n = 256.

samples: total=1558926 fraction of samples 0.1 0.01 0.001 0.0001 105.00 5.00 25.00 45.00 65.00 85.00 overhead in microseconds (bin size = 10.00us)

Fig. 1718: Histogram of DSP handler overhead under G-EDF-MP, with m = 64 and n = 256.













min=0.07us max=111.59us avg=3.59us median=3.25us stdev=2.56us



Fig. 1722: Histogram of release latency under SD, with m = 64 and n = 256.



Fig. 1723: Histogram of tick overhead under SD, with m = 64 and n = 256.



Fig. 1724: Histogram of context-switching overhead under SD, with m = 64 and n = 256.



Fig. 1725: Histogram of IPI latency under SD, with m = 64 and n = 256.

min=0.20us max=1264.63us avg=244.87us median=200.85us stdev=192.89us



Fig. 1726: Histogram of scheduling overhead (a) under GSN-EDF, with m=64 and n=320.



Fig. 1727: Histogram of scheduling overhead (b) under GSN-EDF, with m = 64 and n = 320.



Fig. 1728: Histogram of release overhead under GSN-EDF, with m = 64 and n = 320.



Fig. 1729: Histogram of release latency under GSN-EDF, with m=64 and n=320.



Fig. 1730: Histogram of tick overhead under GSN-EDF, with m = 64and n = 320.



Fig. 1731: Histogram of context-switching overhead under GSN-EDF, with m = 64 and n = 320.



Fig. 1732: Histogram of IPI latency under GSN-EDF, with m = 64 and n = 320.



Fig. 1733: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 64 and n = 320.

samples: total=2053254 fraction of samples 0.1 0.01 0.001 0.0001 0.30 1.50 2.70 3.90 5.10 6.30 7.50 8.70 9.90



min=1.95us max=739.91us avg=169.65us median=122.76us stdev=136.91us



Fig. 1735: Histogram of release overhead under GSN-EDF-DI, with m = 64 and n = 320.









overhead in microseconds (bin size = 0.60us)

230





Fig. 1738: Histogram of context-switching overhead under GSN-EDF-DI, with m = 64 and n = 320.



Fig. 1739: Histogram of IPI latency under GSN-EDF-DI, with m = 64and n = 320.



Fig. 1740: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 64 and n = 320.



Fig. 1741: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 64 and n = 320.

min=0.79us max=230.87us avg=74.12us median=67.67us stdev=43.04us samples: total=97429 fraction of samples 0. 0.01 0.00 0.0001 7.50 37.50 67.50 97.50 127.50 157.50 187.50 217.50

overhead in microseconds (bin size = 15.00us) Fig. 1742: Histogram of release overhead under G-EDF-MP, with m =64 and n = 320.

min=8.38us max=113.91us avg=36.61us median=33.65us stdev=14.20us



Fig. 1743: Histogram of release latency under G-EDF-MP, with m = 64and n = 320.



Fig. 1744: Histogram of tick overhead under G-EDF-MP, with m = 64and n = 320.



Fig. 1745: Histogram of context-switching overhead under G-EDF-MP, with m = 64 and n = 320.



Fig. 1746: Histogram of IPI latency under G-EDF-MP, with m = 64 and n = 320.



Fig. 1747: Histogram of client request latency under G-EDF-MP, with m = 64 and n = 320.



Fig. 1748: Histogram of DSP handler overhead under G-EDF-MP, with m = 64 and n = 320.



Fig. 1749: Histogram of scheduling overhead (a) under SD, with m = 64 and n = 320.

min=0.06us max=896.98us avg=10.15us median=0.08us stdev=29.56us





min=0.20us max=991.80us avg=51.96us median=25.15us stdev=68.58us



Fig. 1751: Histogram of release overhead under SD, with m = 64 and n = 320.











Fig. 1754: Histogram of context-switching overhead under SD, with m = 64 and n = 320.



Fig. 1755: Histogram of IPI latency under SD, with m = 64 and n = 320.



Fig. 1756: Histogram of scheduling overhead (a) under GSN-EDF, with m = 64 and n = 384.



Fig. 1757: Histogram of scheduling overhead (b) under GSN-EDF, with m=64 and n=384.

min=1.45us max=945.05us avg=244.72us median=184.47us stdev=189.89us



overhead in microseconds (bin size = 100.00us)

Fig. 1758: Histogram of release overhead under GSN-EDF, with m = 64 and n = 384.









min=0.36us max=16.38us avg=2.00us median=1.89us stdev=0.78us samples: total=2526119 fraction of samples 0. 0.01 0.001 0.0001 0.50 2.50 4.50 6.50 8.50 10.50 12.50 14.50 16.50 overhead in microseconds (bin size = 1.00us)





Fig. 1762: Histogram of IPI latency under GSN-EDF, with m = 64 and n = 384.



Fig. 1763: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 64 and n = 384.



Fig. 1764: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 64 and n = 384.



Fig. 1765: Histogram of release overhead under GSN-EDF-DI, with m=64 and n=384.

min=4.75us max=104.61us avg=44.59us median=42.45us stdev=13.45us samples: total=98465 0.001 0.0001 5.00 25.00 45.00 65.00 85.00 105.00

overhead in microseconds (bin size = 10.00us) Fig. 1766: Histogram of release latency under GSN-EDF-DI, with m = 64 and n = 384.







Fig. 1768: Histogram of context-switching overhead under GSN-EDF-DI, with m = 64 and n = 384.



Fig. 1769: Histogram of IPI latency under GSN-EDF-DI, with m=64 and n=384.



Fig. 1770: Histogram of scheduling overhead (a) under G-EDF-MP, with m=64 and n=384.



Fig. 1771: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 64 and n = 384.



Fig. 1772: Histogram of release overhead under G-EDF-MP, with m = 64 and n = 384.



Fig. 1773: Histogram of release latency under G-EDF-MP, with m=64 and n=384.

min=0.15us max=21.82us avg=0.94us median=0.25us stdev=2.07us



Fig. 1774: Histogram of tick overhead under G-EDF-MP, with m=64 and n=384.















Fig. 1778: Histogram of DSP handler overhead under G-EDF-MP, with m=64 and n=384.



Fig. 1779: Histogram of scheduling overhead (a) under SD, with m = 64 and n = 384.



Fig. 1780: Histogram of scheduling overhead (b) under SD, with m = 64 and n = 384.



Fig. 1781: Histogram of release overhead under SD, with m = 64 and n = 384.

min=1.48us max=2005.69us avg=45.92us median=14.96us stdev=96.92us 0.01 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001 0.001

Fig. 1782: Histogram of release latency under SD, with m = 64 and n = 384.



Fig. 1783: Histogram of tick overhead under SD, with m = 64 and n = 384.



Fig. 1784: Histogram of context-switching overhead under SD, with m = 64 and n = 384.



Fig. 1785: Histogram of IPI latency under SD, with m = 64 and n = 384.



Fig. 1786: Histogram of scheduling overhead (a) under GSN-EDF, with m = 64 and n = 448.



Fig. 1787: Histogram of scheduling overhead (b) under GSN-EDF, with m = 64 and n = 448.



Fig. 1788: Histogram of release overhead under GSN-EDF, with m =64 and n = 448.



Fig. 1789: Histogram of release latency under GSN-EDF, with m = 64and n = 448.

min=0.15us max=163.79us avg=1.25us median=0.43us stdev=2.30us samples: total=5914797 fraction of samples 0.1 0.01 0.001 0.0001 5.00 25.00 45.00 65.00 85.00 105.00 125.00 145.00 165.00 overhead in microseconds (bin size = 10.00us)

Fig. 1790: Histogram of tick overhead under GSN-EDF, with m = 64and n = 448.



Fig. 1791: Histogram of context-switching overhead under GSN-EDF, with m = 64 and n = 448.











Fig. 1794: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 64 and n = 448.



Fig. 1795: Histogram of release overhead under GSN-EDF-DI, with m = 64 and n = 448.



Fig. 1796: Histogram of release latency under GSN-EDF-DI, with m =64 and n = 448.



Fig. 1797: Histogram of tick overhead under GSN-EDF-DI, with m =64 and n = 448.

min=0.35us max=21.44us avg=2.08us median=1.97us stdev=0.81us samples: total=2911853 fraction of samples 0.1 0.01 0.001 0.0001 1.00 5.00 9.00 13.00 17.00 21.00 overhead in microseconds (bin size = 2.00us)

Fig. 1798: Histogram of context-switching overhead under GSN-EDF-DI, with m = 64 and n = 448.







Fig. 1800: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 64 and n = 448.







Fig. 1802: Histogram of release overhead under G-EDF-MP, with m = 64 and n = 448.



Fig. 1803: Histogram of release latency under G-EDF-MP, with m = 64 and n = 448.



Fig. 1804: Histogram of tick overhead under G-EDF-MP, with m = 64 and n = 448.



Fig. 1805: Histogram of context-switching overhead under G-EDF-MP, with m = 64 and n = 448.

min=0.51us max=79.18us avg=1.97us median=1.69us stdev=1.32us samples: total=2080913 fraction of samples 0.1 0.01 0.001 0.0001 2.50 12.50 22.50 32.50 42.50 52.50 62.50 72.50 overhead in microseconds (bin size = 5.00us)





overhead in microseconds (bin size = 20.00us) Fig. 1807: Histogram of client request latency under G-EDF-MP, with m = 64 and n = 448.

170.00

210.00

250.00

290.00

130.00









10.00

50.00

90.00





Fig. 1810: Histogram of scheduling overhead (b) under SD, with m = 64 and n = 448.



Fig. 1811: Histogram of release overhead under SD, with m = 64 and n = 448.



Fig. 1812: Histogram of release latency under SD, with m = 64 and n = 448.



Fig. 1813: Histogram of tick overhead under SD, with m = 64 and n = 448.

min=0.21us max=30.68us avg=1.78us median=1.69us stdev=1.12us samples: total=3201476 fraction of samples 0.1 0.01 0.001 0.0001 5.00 1.00 9.00 13.00 17.00 21.00 25.00 29.00 overhead in microseconds (bin size = 2.00us)

Fig. 1814: Histogram of context-switching overhead under SD, with m = 64 and n = 448.



Fig. 1815: Histogram of IPI latency under SD, with m = 64 and n = 448.











Fig. 1818: Histogram of release overhead under GSN-EDF, with m = 64 and n = 512.



Fig. 1819: Histogram of release latency under GSN-EDF, with m = 64 and n = 512.



Fig. 1820: Histogram of tick overhead under GSN-EDF, with m = 64 and n = 512.



Fig. 1821: Histogram of context-switching overhead under GSN-EDF, with m = 64 and n = 512.



Fig. 1822: Histogram of IPI latency under GSN-EDF, with m = 64 and n = 512.





Fig. 1823: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m = 64 and n = 512.







Fig. 1825: Histogram of release overhead under GSN-EDF-DI, with m=64 and n=512.



Fig. 1826: Histogram of release latency under GSN-EDF-DI, with m = 64 and n = 512.



Fig. 1827: Histogram of tick overhead under GSN-EDF-DI, with m = 64 and n = 512.



Fig. 1828: Histogram of context-switching overhead under GSN-EDF-DI, with m = 64 and n = 512.



Fig. 1829: Histogram of IPI latency under GSN-EDF-DI, with m = 64 and n = 512.

min=0.16us max=105.55us avg=2.05us median=1.77us stdev=1.24us amples: total=3522029 0.01 0.001 5.00 25.00 45.00 65.00 85.00 105.00 overhead in microseconds (bin size = 10.00us)

Fig. 1830: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 64 and n = 512.







Fig. 1832: Histogram of release overhead under G-EDF-MP, with m = 64 and n = 512.



Fig. 1833: Histogram of release latency under G-EDF-MP, with m=64 and n=512.



Fig. 1834: Histogram of tick overhead under G-EDF-MP, with m = 64and n = 512.



Fig. 1835: Histogram of context-switching overhead under G-EDF-MP, with m = 64 and n = 512.



Fig. 1836: Histogram of IPI latency under G-EDF-MP, with m = 64and n = 512.



Fig. 1837: Histogram of client request latency under G-EDF-MP, with m = 64 and n = 512.

min=0.07us max=113.97us avg=4.28us median=3.32us stdev=2.91us samples: total=3165039 fraction of samples 0.1 0.01 0.001 0.0001 105.00 85.00









Fig. 1840: Histogram of scheduling overhead (b) under SD, with m =64 and n = 512.





5.00 25.00 45.00 65.00

m = 64 and n = 512.



overhead in microseconds (bin size = 200.00us) Fig. 1842: Histogram of release latency under SD, with m=64 and n=512.



Fig. 1843: Histogram of tick overhead under SD, with m = 64 and n = 512.



Fig. 1844: Histogram of context-switching overhead under SD, with m = 64 and n = 512.



Fig. 1845: Histogram of IPI latency under SD, with m = 64 and n = 512.

min=0.26us max=2414.83us avg=480.45us median=455.25us stdev=297.00us



Fig. 1846: Histogram of scheduling overhead (a) under GSN-EDF, with m=64 and n=576.







Fig. 1848: Histogram of release overhead under GSN-EDF, with m = 64 and n = 576.



Fig. 1849: Histogram of release latency under GSN-EDF, with m = 64 and n = 576.



Fig. 1850: Histogram of tick overhead under GSN-EDF, with m = 64 and n = 576.



Fig. 1851: Histogram of context-switching overhead under GSN-EDF, with m = 64 and n = 576.



Fig. 1852: Histogram of IPI latency under GSN-EDF, with m = 64 and n = 576.



Fig. 1853: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m=64 and n=576.

samples: total=4083006 fraction of samples 0.1 0.01 0.001 0.0001 0.45 2.25 4.05 5.85 7.65 9.45 11.25 13.05 14.85 overhead in microseconds (bin size = 0.90us)



min=1.40us max=1444.26us avg=360.74us median=335.91us stdev=227.94us



Fig. 1855: Histogram of release overhead under GSN-EDF-DI, with m=64 and n=576.



Fig. 1856: Histogram of release latency under GSN-EDF-DI, with m = 64 and n = 576.



Fig. 1857: Histogram of tick overhead under GSN-EDF-DI, with m = 64 and n = 576.

min=0.07us max=14.78us avg=0.10us median=0.07us stdev=0.13us



Fig. 1858: Histogram of context-switching overhead under GSN-EDF-DI, with m = 64 and n = 576.



Fig. 1859: Histogram of IPI latency under GSN-EDF-DI, with m = 64 and n = 576.



Fig. 1860: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 64 and n = 576.



Fig. 1861: Histogram of scheduling overhead (b) under G-EDF-MP, with m=64 and n=576.

overhead in microseconds (bin size = 15.00us) Fig. 1862: Histogram of release overhead under G-EDF-MP, with m = 64 and n = 576.

min=8.96us max=115.36us avg=36.80us median=33.32us stdev=16.16us



Fig. 1863: Histogram of release latency under G-EDF-MP, with m = 64 and n = 576.



Fig. 1864: Histogram of tick overhead under G-EDF-MP, with m = 64 and n = 576.



Fig. 1865: Histogram of context-switching overhead under G-EDF-MP, with m=64 and n=576.



Fig. 1866: Histogram of IPI latency under G-EDF-MP, with m = 64 and n = 576.



Fig. 1867: Histogram of client request latency under G-EDF-MP, with m = 64 and n = 576.



Fig. 1868: Histogram of DSP handler overhead under G-EDF-MP, with m = 64 and n = 576.



Fig. 1869: Histogram of scheduling overhead (a) under SD, with m = 64 and n = 576.

min=0.06us max=942.47us avg=13.88us median=0.08us stdev=43.03us



Fig. 1870: Histogram of scheduling overhead (b) under SD, with m = 64 and n = 576.



Fig. 1871: Histogram of release overhead under SD, with m = 64 and n = 576.







Fig. 1873: Histogram of tick overhead under SD, with m = 64 and n = 576.



Fig. 1874: Histogram of context-switching overhead under SD, with m = 64 and n = 576.



Fig. 1875: Histogram of IPI latency under SD, with m = 64 and n = 576.



Fig. 1876: Histogram of scheduling overhead (a) under GSN-EDF, with m=64 and n=640.



Fig. 1877: Histogram of scheduling overhead (b) under GSN-EDF, with m=64 and n=640.

min=1.23us max=1639.09us avg=477.22us median=468.78us stdev=275.88us 0.1 0.01 0.001 0.001

overhead in microseconds (bin size = 100.00us) Fig. 1878: Histogram of release overhead under GSN-EDF, with m = 64 and n = 640.

850.0

1050.0

1250.0

1450.0

1650.0

650.0

50.0

250.0

450.0

min=1.79us max=1714.68us avg=80.87us median=36.51us stdev=143.80us



Fig. 1879: Histogram of release latency under GSN-EDF, with m = 64 and n = 640.







Fig. 1881: Histogram of context-switching overhead under GSN-EDF, with m=64 and n=640.



Fig. 1882: Histogram of IPI latency under GSN-EDF, with m=64 and n=640.



Fig. 1883: Histogram of scheduling overhead (a) under GSN-EDF-DI, with m=64 and n=640.



Fig. 1884: Histogram of scheduling overhead (b) under GSN-EDF-DI, with m = 64 and n = 640.



Fig. 1885: Histogram of release overhead under GSN-EDF-DI, with m=64 and n=640.

min=1.65us max=30327.49us avg=133.11us median=37.34us stdev=1131.39us



Fig. 1886: Histogram of release latency under GSN-EDF-DI, with m = 64 and n = 640.







Fig. 1888: Histogram of context-switching overhead under GSN-EDF-DI, with m=64 and n=640.







Fig. 1890: Histogram of scheduling overhead (a) under G-EDF-MP, with m = 64 and n = 640.



Fig. 1891: Histogram of scheduling overhead (b) under G-EDF-MP, with m = 64 and n = 640.



Fig. 1892: Histogram of release overhead under G-EDF-MP, with m = 64 and n = 640.



Fig. 1893: Histogram of release latency under G-EDF-MP, with m=64 and n=640.

min=0.15us max=19.96us avg=1.17us median=0.28us stdev=2.29us



Fig. 1894: Histogram of tick overhead under G-EDF-MP, with m = 64 and n = 640.















Fig. 1898: Histogram of DSP handler overhead under G-EDF-MP, with m=64 and n=640.



Fig. 1899: Histogram of scheduling overhead (a) under SD, with m = 64 and n = 640.



Fig. 1900: Histogram of scheduling overhead (b) under SD, with m = 64 and n = 640.



Fig. 1901: Histogram of release overhead under SD, with m=64 and n=640.

min=1.30us max=2199.03us avg=71.09us median=18.55us stdev=145.55us



Fig. 1902: Histogram of release latency under SD, with m = 64 and n = 640.







Fig. 1904: Histogram of context-switching overhead under SD, with m = 64 and n = 640.



Fig. 1905: Histogram of IPI latency under SD, with m = 64 and n = 640.

E. All Schedulability Results

In the following, all schedulability results are reported for the sake of completeness and transparency. As discussed in Sec. V, we used two different task set generation methods to assess schedulability both as a function of total utilization and as a function of the number of tasks.

We begin with the results that show hard and soft realtime schedulability as a function of increasing total utilization, which we obtained using Emberson *et al.*'s task set generation method [24]. The graphs are organized as follows.

- Figs. 1906–1941 show results for m = 8.
 - Figs. 1906–1909 show results for n = 2m = 16.
 - Figs. 1910–1913 show results for n = 3m = 24.
 - Figs. 1914–1917 show results for n = 4m = 32.
 - Figs. 1918–1921 show results for n = 5m = 40.
 - Figs. 1922–1925 show results for n = 6m = 48.
 - Figs. 1926–1929 show results for n = 7m = 56.
 - Figs. 1930–1933 show results for n = 8m = 64.
 - Figs. 1934–1937 show results for n = 9m = 72.
 - Figs. 1938–1941 show results for n = 10m = 80.
- Figs. 1942–1977 show results for m = 16.
 - Figs. 1942–1945 show results for n = 2m = 32.
 - Figs. 1946–1949 show results for n = 3m = 48.
 - Figs. 1950–1953 show results for n = 4m = 64.
 - Figs. 1954–1957 show results for n = 5m = 80.
 - Figs. 1958–1961 show results for n = 6m = 96.
 - Figs. 1962–1965 show results for n = 7m = 112.
 - Figs. 1966–1969 show results for n = 8m = 128.
 - Figs. 1970–1973 show results for n = 9m = 144.
 - Figs. 1974–1977 show results for n = 10m = 160.
- Figs. 1978–2013 show results for m = 24.
 - Figs. 1978–1981 show results for n = 2m = 48.
 - Figs. 1982–1985 show results for n = 3m = 72.
 - Figs. 1986–1989 show results for n = 4m = 96.
 - Figs. 1990–1993 show results for n = 5m = 120.
 - Figs. 1994–1997 show results for n = 6m = 144.
 - Figs. 1998–2001 show results for n = 7m = 168.
 - Figs. 2002–2005 show results for n = 8m = 192.
 - Figs. 2006–2009 show results for n = 9m = 216.
 - Figs. 2010–2013 show results for n = 10m = 240.
- Figs. 2014–2049 show results for m = 32.
 - Figs. 2014–2017 show results for n = 2m = 64.
 - Figs. 2018–2021 show results for n = 3m = 96.
 - Figs. 2022–2025 show results for n = 4m = 128.
 - Figs. 2026–2029 show results for n = 5m = 160.
 - Figs. 2030–2033 show results for n = 6m = 192.
 - Figs. 2034–2037 show results for n = 7m = 224.
 - Figs. 2038–2041 show results for n = 8m = 256.
 - Figs. 2042–2045 show results for n = 9m = 288.
 - Figs. 2046–2049 show results for n = 10m = 320.
- Figs. 2050–2085 show results for m = 48.
 - Figs. 2050–2053 show results for n = 2m = 96.
 - Figs. 2054–2057 show results for n = 3m = 144.
 - Figs. 2058–2061 show results for n = 4m = 192.
 - Figs. 2062–2065 show results for n = 5m = 240.

- Figs. 2066–2069 show results for n = 6m = 288.
- Figs. 2070–2073 show results for n = 7m = 336.
- Figs. 2074–2077 show results for n = 8m = 384.
- Figs. 2078–2081 show results for n = 9m = 432.
- Figs. 2082–2085 show results for n = 10m = 480.
- Figs. 2086–2120 show results for m = 64.
 - Figs. 2086–2089 show results for n = 2m = 128.
 - Figs. 2090–2093 show results for n = 3m = 192.
 - Figs. 2094–2097 show results for n = 4m = 256.
 - Figs. 2098–2101 show results for n = 5m = 320.
 - Figs. 2102–2105 show results for n = 6m = 384.
 - Figs. 2106–2109 show results for n = 7m = 448.
 - Figs. 2110–2112 show results for n = 8m = 512.
 - Figs. 2113–2116 show results for n = 9m = 576.
 - Figs. 2117–2120 show results for n = 10m = 640.

In the second set of schedulability experiments, we did not generate task sets for a fixed total utilization. Instead, we generated an increasing number of tasks by drawing per-task utilizations and periods from various distributions (adopted from the schedulability experiments in [13, 16, 18]) and let the total utilization vary as a result. The resulting schedulability plots are organized as follows.

- Figs. 2121–2264 show results for m = 8.
 - Figs. 2121–2144 show results with periods ranging from 1 ms to 1000 ms.
 - Figs. 2145–2168 show results with periods ranging from 3 ms to 33 ms.
 - Figs. 2169–2192 show results with periods ranging from 10 ms to 100 ms.
 - Figs. 2193–2216 show results with periods ranging from 10 ms to 1000 ms.
 - Figs. 2217–2240 show results with periods ranging from 50 ms to 250 ms.
 - Figs. 2241–2264 show results with periods ranging from 200 ms to 1000 ms.
- Figs. 2265–2408 show results for m = 16.
 - Figs. 2265–2288 show results with periods ranging from 1 ms to 1000 ms.
 - Figs. 2289–2312 show results with periods ranging from 3 ms to 33 ms.
 - Figs. 2313–2336 show results with periods ranging from 10 ms to 100 ms.
 - Figs. 2337–2360 show results with periods ranging from 10 ms to 1000 ms.
 - Figs. 2361–2384 show results with periods ranging from 50 ms to 250 ms.
 - Figs. 2385–2408 show results with periods ranging from 200 ms to 1000 ms.
- Figs. 2409–2552 show results for m = 24.
 - Figs. 2409–2432 show results with periods ranging from 1 ms to 1000 ms.
 - Figs. 2433–2456 show results with periods ranging from 3 ms to 33 ms.
 - Figs. 2457–2480 show results with periods ranging from 10 ms to 100 ms.
- Figs. 2481–2504 show results with periods ranging from 10 ms to 1000 ms.
- Figs. 2505–2528 show results with periods ranging from 50 ms to 250 ms.
- Figs. 2529–2552 show results with periods ranging from 200 ms to 1000 ms.
- Figs. 2553–2696 show results for m = 32.
 - Figs. 2553–2576 show results with periods ranging from 1 ms to 1000 ms.
 - Figs. 2577–2600 show results with periods ranging from 3 ms to 33 ms.
 - Figs. 2601–2624 show results with periods ranging from 10 ms to 100 ms.
 - Figs. 2625–2648 show results with periods ranging from 10 ms to 1000 ms.
 - Figs. 2649–2672 show results with periods ranging from 50 ms to 250 ms.
 - Figs. 2673–2696 show results with periods ranging from 200 ms to 1000 ms.
- Figs. 2697–2839 show results for m = 48.
 - Figs. 2697–2719 show results with periods ranging from 1 ms to 1000 ms.
 - Figs. 2720–2743 show results with periods ranging from 3 *ms* to 33 *ms*.
 - Figs. 2744–2767 show results with periods ranging from 10 ms to 100 ms.
 - Figs. 2768–2791 show results with periods ranging from 10 ms to 1000 ms.
 - Figs. 2792–2815 show results with periods ranging from 50 ms to 250 ms.
 - Figs. 2816–2839 show results with periods ranging from 200 ms to 1000 ms.
- Figs. 2840–2982 show results for m = 64.
 - Figs. 2840–2862 show results with periods ranging from 1 ms to 1000 ms.
 - Figs. 2863–2886 show results with periods ranging from 3 ms to 33 ms.
 - Figs. 2887–2910 show results with periods ranging from 10 ms to 100 ms.
 - Figs. 2911–2934 show results with periods ranging from 10 ms to 1000 ms.
 - Figs. 2935–2958 show results with periods ranging from 50 ms to 250 ms.
 - Figs. 2959–2982 show results with periods ranging from 200 ms to 1000 ms.



Fig. 1906: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 2m = 16, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1907: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 2m = 16, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1908: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 2m = 16, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1909: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 2m = 16, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1910: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 3m = 24, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1911: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 3m = 24, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1912: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 3m = 24, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1913: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 3m = 24, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1914: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 4m = 32, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1915: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 4m = 32, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1916: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 4m = 32, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1917: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 4m = 32, and log-uniformly distributed periods in [1ms, 1000ms].



ratio of schedulable task sets

Fig. 1918: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 5m = 40, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1919: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 5m = 40, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1920: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 5m = 40, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1921: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 5m = 40, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1922: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 6m = 48, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1923: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 6m = 48, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1924: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 6m = 48, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1925: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 6m = 48, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1926: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 7m = 56, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1927: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 7m = 56, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1928: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 7m = 56, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1929: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 7m = 56, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1930: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 8m = 64, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1931: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 8m = 64, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1932: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 8m = 64, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1933: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 8m = 64, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1934: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 9m = 72, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1935: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 9m = 72, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1936: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 9m = 72, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1937: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 9m = 72, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1938: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 10m = 80, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1939: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 10m = 80, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1940: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 10m = 80, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1941: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 8, n = 10m = 80, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1942: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 2m = 32, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1943: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 2m = 32, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1944: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 2m = 32, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1945: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 2m = 32, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1946: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 3m = 48, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1947: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 3m = 48, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1948: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 3m = 48, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1949: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 3m = 48, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1950: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 4m = 64, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1951: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 4m = 64, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1952: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 4m = 64, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1953: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 4m = 64, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1954: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 5m = 80, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1955: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 5m = 80, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1956: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 5m = 80, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1957: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 5m = 80, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1958: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 6m = 96, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1959: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 6m = 96, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1960: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 6m = 96, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1961: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 6m = 96, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1962: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 7m = 112, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1963: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 7m = 112, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1964: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 7m = 112, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1965: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 7m = 112, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1966: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 8m = 128, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1967: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 8m = 128, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1968: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 8m = 128, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1969: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 8m = 128, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1970: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 9m = 144, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1971: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 9m = 144, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1972: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 9m = 144, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1973: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 9m = 144, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1974: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 10m = 160, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1975: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 10m = 160, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1976: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 10m = 160, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1977: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 16, n = 10m = 160, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1978: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 2m = 48, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1979: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 2m = 48, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1980: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 2m = 48, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1981: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 2m = 48, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1982: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 3m = 72, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1983: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 3m = 72, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1984: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 3m = 72, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1985: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 3m = 72, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1986: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 4m = 96, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1987: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 4m = 96, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1988: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 4m = 96, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1989: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 4m = 96, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1990: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 5m = 120, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1991: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 5m = 120, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1992: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 5m = 120, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1993: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 5m = 120, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1994: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 6m = 144, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1995: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 6m = 144, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1996: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 6m = 144, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1997: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 6m = 144, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 1998: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 7m = 168, and uniformly distributed periods in [1ms, 1000ms].



Fig. 1999: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 7m = 168, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2000: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 7m = 168, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2001: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 7m = 168, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2002: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 8m = 192, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2003: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 8m = 192, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2004: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 8m = 192, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2005: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 8m = 192, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2006: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 9m = 216, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2007: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 9m = 216, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2008: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 9m = 216, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2009: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 9m = 216, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2010: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 10m = 240, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2011: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 10m = 240, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2012: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 10m = 240, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2013: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 24, n = 10m = 240, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2014: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 2m = 64, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2015: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 2m = 64, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2016: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 2m = 64, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2017: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 2m = 64, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2018: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 3m = 96, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2019: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 3m = 96, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2020: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 3m = 96, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2021: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 3m = 96, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2022: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 4m = 128, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2023: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 4m = 128, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2024: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 4m = 128, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2025: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 4m = 128, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2026: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 5m = 160, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2027: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 5m = 160, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2028: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 5m = 160, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2029: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 5m = 160, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2030: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 6m = 192, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2031: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 6m = 192, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2032: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 6m = 192, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2033: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 6m = 192, and log-uniformly distributed periods in [1ms, 1000ms].



star open point of the second second

Fig. 2034: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 7m = 224, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2035: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 7m = 224, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2036: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 7m = 224, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2037: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 7m = 224, and log-uniformly distributed periods in [1ms, 1000ms].

Fig. 2038: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 8m = 256, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2039: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 8m = 256, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2040: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 8m = 256, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2041: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 8m = 256, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2042: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 9m = 288, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2043: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 9m = 288, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2044: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 9m = 288, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2045: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 9m = 288, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2046: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 10m = 320, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2047: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 10m = 320, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2048: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 10m = 320, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2049: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 32, n = 10m = 320, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2050: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 2m = 96, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2051: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 2m = 96, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2052: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 2m = 96, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2053: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 2m = 96, and log-uniformly distributed periods in [1ms, 1000ms].



ratio of schedulable task sets

Fig. 2054: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 3m = 144, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2055: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 3m = 144, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2056: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 3m = 144, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2057: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 3m = 144, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2058: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 4m = 192, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2059: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 4m = 192, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2060: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 4m = 192, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2061: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 4m = 192, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2062: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 5m = 240, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2063: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 5m = 240, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2064: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 5m = 240, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2065: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 5m = 240, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2066: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 6m = 288, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2067: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 6m = 288, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2068: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 6m = 288, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2069: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 6m = 288, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2070: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 7m = 336, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2071: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 7m = 336, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2072: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 7m = 336, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2073: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 7m = 336, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2074: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 8m = 384, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2075: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 8m = 384, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2076: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 8m = 384, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2077: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 8m = 384, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2078: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 9m = 432, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2079: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 9m = 432, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2080: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 9m = 432, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2081: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 48, n = 9m = 432, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2082: Hard schedulability (i.e., no deadline misses allowed) as a function of total utilization using Emberson et al.'s task set generation method [24] with m = 48, n = 10m = 480, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2083: Hard schedulability (i.e., no deadline misses allowed) as a function of total utilization using Emberson et al.'s task set generation method [24] with m = 48, n = 10m = 480, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2084: Soft schedulability (i.e., bounded tardiness allowed) as a function of total utilization using Emberson et al.'s task set generation method [24] with m = 48, n = 10m = 480, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2085: Soft schedulability (i.e., bounded tardiness allowed) as a function of total utilization using Emberson et al.'s task set generation method [24] with m = 48, n = 10m = 480, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2086: Hard schedulability (i.e., no deadline misses allowed) as a function of total utilization using Emberson et al.'s task set generation method [24] with m = 64, n = 2m = 128, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2087: Hard schedulability (i.e., no deadline misses allowed) as a function of total utilization using Emberson et al.'s task set generation method [24] with m = 64, n = 2m = 128, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2088: Soft schedulability (i.e., bounded tardiness allowed) as a function of total utilization using Emberson et al.'s task set generation method [24] with m = 64, n = 2m = 128, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2089: Soft schedulability (i.e., bounded tardiness allowed) as a function of total utilization using Emberson et al.'s task set generation method [24] with m = 64, n = 2m = 128, and log-uniformly distributed periods in [1ms, 1000ms].

sets



Fig. 2090: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 3m = 192, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2091: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 3m = 192, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2092: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 3m = 192, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2093: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 3m = 192, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2094: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 4m = 256, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2095: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 4m = 256, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2096: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 4m = 256, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2097: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 4m = 256, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2098: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 5m = 320, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2099: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 5m = 320, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2100: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 5m = 320, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2101: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 5m = 320, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2102: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 6m = 384, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2103: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 6m = 384, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2104: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 6m = 384, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2105: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 6m = 384, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2106: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 7m = 448, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2107: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 7m = 448, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2108: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 7m = 448, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2109: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 7m = 448, and log-uniformly distributed periods in [1ms, 1000ms].



ratio of schedulable task sets

Fig. 2110: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 8m = 512, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2111: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 8m = 512, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2112: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 8m = 512, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2113: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 9m = 576, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2114: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 9m = 576, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2115: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 9m = 576, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2116: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 9m = 576, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2117: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 10m = 640, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2118: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 10m = 640, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2119: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 10m = 640, and uniformly distributed periods in [1ms, 1000ms].



Fig. 2120: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of total utilization using Emberson *et al.*'s task set generation method [24] with m = 64, n = 10m = 640, and log-uniformly distributed periods in [1ms, 1000ms].



Fig. 2121: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2122: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2123: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2124: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2125: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2126: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2127: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2128: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2129: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2130: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2131: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2132: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2133: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2134: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2135: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2136: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2137: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2138: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2139: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2140: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2141: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2142: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2143: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2144: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2145: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [3ms, 33ms].



Fig. 2146: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2147: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [3ms, 33ms].



Fig. 2148: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2149: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [3ms, 33ms].



Fig. 2150: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2151: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [3ms, 33ms].



Fig. 2152: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2153: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [3ms, 33ms].



Fig. 2154: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2155: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [3ms, 33ms].



Fig. 2156: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2157: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [3ms, 33ms].



Fig. 2158: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2159: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [3ms, 33ms].



Fig. 2160: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2161: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [3ms, 33ms].



Fig. 2162: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2163: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [3ms, 33ms].



Fig. 2164: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2165: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [3ms, 33ms].



Fig. 2166: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2167: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [3ms, 33ms].



Fig. 2168: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2169: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 100ms].



Fig. 2170: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2171: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 100ms].



Fig. 2172: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2173: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 100ms].



Fig. 2174: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2175: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 100ms].



Fig. 2176: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2177: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 100ms].



Fig. 2178: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2179: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 100ms].



Fig. 2180: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2181: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 100ms].



Fig. 2182: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2183: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 100ms].



Fig. 2184: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 100ms].


Fig. 2185: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 100ms].



Fig. 2186: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2187: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 100ms].



Fig. 2188: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2189: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 100ms].



Fig. 2190: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2191: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 100ms].



Fig. 2192: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2193: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2194: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2195: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2196: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2197: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2198: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2199: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2200: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2201: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2202: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2203: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2204: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2205: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2206: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2207: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2208: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2209: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2210: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2211: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2212: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2213: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2214: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2215: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2216: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2217: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [50ms, 250ms].



Fig. 2218: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2219: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [50ms, 250ms].



Fig. 2220: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2221: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [50ms, 250ms].



Fig. 2222: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2223: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [50ms, 250ms].



Fig. 2224: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2225: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [50ms, 250ms].



Fig. 2226: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2227: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [50ms, 250ms].



Fig. 2228: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2229: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [50ms, 250ms].



Fig. 2230: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2231: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [50ms, 250ms].



Fig. 2232: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2233: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [50ms, 250ms].



Fig. 2234: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2235: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [50ms, 250ms].



Fig. 2236: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2237: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [50ms, 250ms].



Fig. 2238: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2239: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [50ms, 250ms].



Fig. 2240: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2241: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [200 ms, 1000 ms].



Fig. 2242: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2243: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2244: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2245: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2246: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2247: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2248: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2249: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2250: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2251: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [200*ms*, 1000*ms*].



Fig. 2252: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2253: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2254: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2255: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2256: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2257: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2258: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2259: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2260: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2261: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2262: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2263: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2264: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 8, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2265: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2266: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2267: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2268: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2269: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2270: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2271: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2272: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2273: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2274: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2275: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2276: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2277: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2278: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2279: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2280: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2281: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2282: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2283: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2284: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2285: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2286: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2287: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2288: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2289: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [3ms, 33ms].



Fig. 2290: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2291: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [3ms, 33ms].



Fig. 2292: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2293: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [3ms, 33ms].



Fig. 2294: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2295: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [3ms, 33ms].



Fig. 2296: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2297: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [3ms, 33ms].



Fig. 2298: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2299: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [3ms, 33ms].



Fig. 2300: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2301: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [3ms, 33ms].



Fig. 2302: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2303: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [3ms, 33ms].



Fig. 2304: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2305: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [3ms, 33ms].



Fig. 2306: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2307: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [3ms, 33ms].



Fig. 2308: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2309: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [3ms, 33ms].



Fig. 2310: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2311: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [3ms, 33ms].



Fig. 2312: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2313: Hard schedulability (i.e., no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 100ms].



Fig. 2314: Hard schedulability (i.e., no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2315: Soft schedulability (i.e., bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 100ms].



Fig. 2316: Soft schedulability (i.e., bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2317: Hard schedulability (i.e., no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 100ms].



Fig. 2318: Hard schedulability (i.e., no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2319: Soft schedulability (i.e., bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 100ms].



Fig. 2320: Soft schedulability (i.e., bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 100ms].

sets



Fig. 2321: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 100ms].



Fig. 2322: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2323: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 100ms].



Fig. 2324: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2325: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 100ms].



Fig. 2326: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2327: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 100ms].



Fig. 2328: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2329: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 100ms].



Fig. 2330: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2331: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 100ms].



Fig. 2332: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2333: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 100ms].



Fig. 2334: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2335: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 100ms].



Fig. 2336: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2337: Hard schedulability (i.e., no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2338: Hard schedulability (i.e., no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2339: Soft schedulability (i.e., bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2340: Soft schedulability (i.e., bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2341: Hard schedulability (i.e., no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2342: Hard schedulability (i.e., no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2343: Soft schedulability (i.e., bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2344: Soft schedulability (i.e., bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 1000ms].

sets



Fig. 2345: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2346: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2347: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2348: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2349: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2350: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2351: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2352: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2353: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2354: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2355: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2356: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2357: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2358: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2359: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2360: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2361: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [50ms, 250ms].



Fig. 2362: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2363: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [50ms, 250ms].



Fig. 2364: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2365: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [50ms, 250ms].



Fig. 2366: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2367: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [50ms, 250ms].



Fig. 2368: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [50ms, 250ms].

sets



Fig. 2369: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [50ms, 250ms].



Fig. 2370: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2371: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [50ms, 250ms].



Fig. 2372: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2373: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [50ms, 250ms].



Fig. 2374: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2375: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [50ms, 250ms].



Fig. 2376: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2377: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [50ms, 250ms].



Fig. 2378: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2379: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [50ms, 250ms].



Fig. 2380: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2381: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [50ms, 250ms].



Fig. 2382: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2383: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [50ms, 250ms].



Fig. 2384: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2385: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2386: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2387: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2388: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2389: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2390: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2391: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2392: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2393: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2394: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2395: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2396: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2397: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2398: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2399: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2400: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2401: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2402: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2403: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2404: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2405: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2406: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2407: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2408: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 16, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2409: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2410: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2411: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2412: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2413: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2414: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2415: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2416: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2417: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2418: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2419: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2420: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2421: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2422: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2423: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2424: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2425: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2426: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2427: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2428: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2429: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2430: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2431: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2432: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2433: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [3ms, 33ms].



Fig. 2434: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2435: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [3ms, 33ms].



Fig. 2436: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2437: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [3ms, 33ms].



Fig. 2438: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2439: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [3ms, 33ms].



Fig. 2440: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2441: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [3ms, 33ms].



Fig. 2442: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2443: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [3ms, 33ms].



Fig. 2444: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2445: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [3ms, 33ms].



Fig. 2446: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2447: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [3ms, 33ms].



Fig. 2448: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2449: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [3ms, 33ms].



Fig. 2450: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2451: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [3ms, 33ms].



Fig. 2452: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2453: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [3ms, 33ms].



Fig. 2454: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2455: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [3ms, 33ms].



Fig. 2456: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2457: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 100ms].



Fig. 2458: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2459: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 100ms].



Fig. 2460: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2461: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 100ms].



Fig. 2462: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2463: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 100ms].



Fig. 2464: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2465: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 100ms].



Fig. 2466: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2467: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 100ms].



Fig. 2468: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2469: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 100ms].



Fig. 2470: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2471: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 100ms].



Fig. 2472: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 100ms].


Fig. 2473: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 100ms].



Fig. 2474: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2475: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 100ms].



Fig. 2476: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2477: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 100ms].



Fig. 2478: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2479: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 100ms].



Fig. 2480: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2481: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2482: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2483: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2484: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2485: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2486: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2487: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2488: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2489: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2490: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2491: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2492: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2493: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2494: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2495: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2496: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2497: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2498: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2499: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2500: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2501: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2502: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2503: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2504: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2505: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [50ms, 250ms].



Fig. 2506: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2507: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [50ms, 250ms].



Fig. 2508: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2509: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [50ms, 250ms].



Fig. 2510: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2511: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [50ms, 250ms].



Fig. 2512: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2513: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [50ms, 250ms].



Fig. 2514: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2515: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [50ms, 250ms].



Fig. 2516: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2517: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [50ms, 250ms].



Fig. 2518: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2519: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [50ms, 250ms].



Fig. 2520: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2521: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [50ms, 250ms].



Fig. 2522: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2523: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [50ms, 250ms].



Fig. 2524: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2525: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [50ms, 250ms].



Fig. 2526: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2527: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [50ms, 250ms].



Fig. 2528: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2529: Hard schedulability (i.e., no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2530: Hard schedulability (i.e., no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2531: Soft schedulability (i.e., bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2532: Soft schedulability (i.e., bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2533: Hard schedulability (i.e., no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2534: Hard schedulability (i.e., no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2535: Soft schedulability (i.e., bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2536: Soft schedulability (i.e., bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [200ms, 1000ms].

sets



Fig. 2537: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2538: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2539: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [200*ms*, 1000*ms*].



Fig. 2540: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2541: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2542: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2543: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2544: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2545: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2546: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2547: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2548: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2549: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2550: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2551: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2552: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 24, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2553: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2554: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2555: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2556: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2557: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2558: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2559: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2560: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2561: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2562: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2563: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2564: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2565: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2566: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2567: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2568: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2569: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2570: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2571: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2572: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2573: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2574: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2575: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2576: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2577: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [3ms, 33ms].



Fig. 2578: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2579: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [3ms, 33ms].



Fig. 2580: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2581: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [3ms, 33ms].



Fig. 2582: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2583: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [3ms, 33ms].



Fig. 2584: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2585: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [3ms, 33ms].



Fig. 2586: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2587: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [3ms, 33ms].



Fig. 2588: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2589: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [3ms, 33ms].



Fig. 2590: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2591: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [3ms, 33ms].



Fig. 2592: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2593: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [3ms, 33ms].



Fig. 2594: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2595: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [3ms, 33ms].



Fig. 2596: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [3ms, 33ms].



ratio of schedulable task sets

Fig. 2597: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [3ms, 33ms].



Fig. 2598: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2599: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [3ms, 33ms].



Fig. 2600: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2601: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 100ms].



Fig. 2602: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2603: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 100ms].



Fig. 2604: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2605: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 100ms].



Fig. 2606: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2607: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 100ms].



Fig. 2608: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2609: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 100ms].



Fig. 2610: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2611: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 100ms].



Fig. 2612: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2613: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 100ms].



Fig. 2614: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2615: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 100ms].



Fig. 2616: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2617: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 100ms].



Fig. 2618: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2619: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 100ms].



Fig. 2620: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2621: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 100ms].



Fig. 2622: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2623: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 100ms].



Fig. 2624: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2625: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2626: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2627: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2628: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2629: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2630: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2631: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2632: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2633: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2634: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2635: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2636: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2637: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2638: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2639: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2640: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2641: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2642: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2643: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2644: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2645: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2646: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2647: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2648: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2649: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [50ms, 250ms].



Fig. 2650: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2651: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [50ms, 250ms].



Fig. 2652: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2653: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [50ms, 250ms].



Fig. 2654: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2655: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [50ms, 250ms].



Fig. 2656: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2657: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [50ms, 250ms].



Fig. 2658: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2659: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [50ms, 250ms].



Fig. 2660: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2661: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [50ms, 250ms].



Fig. 2662: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2663: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [50ms, 250ms].



Fig. 2664: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2665: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [50ms, 250ms].



Fig. 2666: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2667: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [50ms, 250ms].



Fig. 2668: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2669: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [50ms, 250ms].



Fig. 2670: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2671: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [50ms, 250ms].



Fig. 2672: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2673: Hard schedulability (i.e., no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2674: Hard schedulability (i.e., no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2675: Soft schedulability (i.e., bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2676: Soft schedulability (i.e., bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2677: Hard schedulability (i.e., no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2678: Hard schedulability (i.e., no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2679: Soft schedulability (i.e., bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2680: Soft schedulability (i.e., bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [200ms, 1000ms].

sets



Fig. 2681: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2682: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2683: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2684: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2685: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2686: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2687: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2688: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2689: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2690: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2691: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2692: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2693: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2694: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2695: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2696: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 32, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2697: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2698: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2699: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2700: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2701: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2702: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2703: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2704: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [1ms, 1000ms].

sets

ratio of schedulable task

ratio of schedulable task sets



Fig. 2705: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2706: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2707: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2708: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2709: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2710: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2711: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2712: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2713: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2714: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2715: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2716: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2717: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2718: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2719: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2720: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [3ms, 33ms].



Fig. 2721: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2722: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [3ms, 33ms].



Fig. 2723: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2724: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [3ms, 33ms].



Fig. 2725: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2726: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [3ms, 33ms].



Fig. 2727: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2728: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [3ms, 33ms].



Fig. 2729: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2730: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [3ms, 33ms].



Fig. 2731: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2732: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [3ms, 33ms].



Fig. 2733: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2734: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [3ms, 33ms].



Fig. 2735: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2736: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [3ms, 33ms].



Fig. 2737: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2738: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [3ms, 33ms].



Fig. 2739: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2740: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [3ms, 33ms].



Fig. 2741: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2742: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [3ms, 33ms].



Fig. 2743: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2744: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 100ms].



Fig. 2745: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2746: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 100ms].



Fig. 2747: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2748: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 100ms].



Fig. 2749: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2750: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 100ms].



Fig. 2751: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2752: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 100ms].



Fig. 2753: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2754: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 100ms].



Fig. 2755: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2756: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 100ms].



Fig. 2757: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2758: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 100ms].



Fig. 2759: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2760: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 100ms].

ratio of schedulable task sets


Fig. 2761: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2762: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 100ms].



Fig. 2763: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2764: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 100ms].



Fig. 2765: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2766: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 100ms].



Fig. 2767: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2768: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2769: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2770: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2771: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2772: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2773: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2774: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2775: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2776: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2777: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2778: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2779: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2780: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2781: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2782: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2783: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2784: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2785: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2786: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2787: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2788: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2789: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2790: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2791: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2792: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [50ms, 250ms].



Fig. 2793: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2794: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [50ms, 250ms].



Fig. 2795: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2796: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [50ms, 250ms].



Fig. 2797: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2798: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [50ms, 250ms].



Fig. 2799: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2800: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [50ms, 250ms].



Fig. 2801: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2802: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [50ms, 250ms].



Fig. 2803: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2804: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [50ms, 250ms].



Fig. 2805: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2806: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [50ms, 250ms].



Fig. 2807: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2808: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [50ms, 250ms].



Fig. 2809: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2810: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [50ms, 250ms].



Fig. 2811: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2812: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [50ms, 250ms].



Fig. 2813: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2814: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [50ms, 250ms].



Fig. 2815: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2816: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2817: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2818: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2819: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2820: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2821: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2822: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2823: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2824: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2825: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2826: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2827: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [200*ms*, 1000*ms*].



Fig. 2828: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2829: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2830: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2831: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2832: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2833: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2834: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2835: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2836: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2837: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2838: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2839: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 48, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2840: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2841: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2842: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2843: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2844: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2845: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2846: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2847: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2848: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2849: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2850: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2851: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2852: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2853: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2854: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2855: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2856: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2857: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [1ms, 1000ms].



Fig. 2858: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2859: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2860: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2861: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [1ms, 1000ms].



Fig. 2862: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [1ms, 1000ms].



Fig. 2863: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [3ms, 33ms].



Fig. 2864: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2865: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [3ms, 33ms].



Fig. 2866: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2867: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [3ms, 33ms].



Fig. 2868: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2869: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [3ms, 33ms].



Fig. 2870: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2871: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [3ms, 33ms].



Fig. 2872: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2873: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [3ms, 33ms].



Fig. 2874: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2875: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [3ms, 33ms].



Fig. 2876: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2877: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [3ms, 33ms].



Fig. 2878: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2879: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [3ms, 33ms].



Fig. 2880: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2881: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [3ms, 33ms].



Fig. 2882: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2883: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [3ms, 33ms].



Fig. 2884: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2885: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [3ms, 33ms].



Fig. 2886: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [3ms, 33ms].



Fig. 2887: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 100ms].



Fig. 2888: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2889: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 100ms].



Fig. 2890: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2891: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 100ms].



Fig. 2892: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2893: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 100ms].



Fig. 2894: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2895: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 100ms].



Fig. 2896: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2897: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 100ms].



Fig. 2898: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2899: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 100ms].



Fig. 2900: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2901: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 100ms].



Fig. 2902: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2903: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 100ms].



Fig. 2904: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2905: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 100ms].



Fig. 2906: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2907: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 100ms].



Fig. 2908: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2909: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 100ms].



Fig. 2910: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 100ms].



Fig. 2911: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2912: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2913: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2914: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2915: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2916: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2917: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2918: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2919: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2920: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2921: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2922: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2923: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2924: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2925: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2926: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2927: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2928: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2929: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [10ms, 1000ms].



Fig. 2930: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2931: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2932: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2933: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [10ms, 1000ms].



Fig. 2934: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [10ms, 1000ms].



Fig. 2935: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [50ms, 250ms].



Fig. 2936: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2937: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [50ms, 250ms].



Fig. 2938: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2939: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [50ms, 250ms].



Fig. 2940: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2941: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [50ms, 250ms].



Fig. 2942: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2943: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [50ms, 250ms].



Fig. 2944: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2945: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [50ms, 250ms].



Fig. 2946: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2947: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [50ms, 250ms].



Fig. 2948: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2949: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [50ms, 250ms].



Fig. 2950: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2951: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [50ms, 250ms].



Fig. 2952: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2953: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [50ms, 250ms].



Fig. 2954: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2955: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [50ms, 250ms].



Fig. 2956: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2957: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [50ms, 250ms].



Fig. 2958: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [50ms, 250ms].



Fig. 2959: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2960: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2961: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2962: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.1, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2963: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2964: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2965: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2966: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.001, 0.1], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2967: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2968: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2969: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2970: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.25, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2971: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2972: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2973: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2974: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.1, 0.4], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2975: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2976: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2977: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods uniformly distributed in [200ms, 1000ms].



Fig. 2978: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations exponentially distributed in (0, 1] with mean 0.5, and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2979: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2980: Hard schedulability (*i.e.*, no deadline misses allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [200ms, 1000ms].



Fig. 2981: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods uniformly distributed in [200ms, 1000ms].



Fig. 2982: Soft schedulability (*i.e.*, bounded tardiness allowed) as a function of n using the task set generation method from [13, 16, 18] for m = 64, per-task utilizations uniformly distributed in [0.5, 0.9], and periods log-uniformly distributed in [200ms, 1000ms].