## Real-Time Multiprocessor Locks with Nesting: Optimizing the Common Case<sup>\*</sup>

Catherine E. Nemitz Tanya Amert James H. Anderson

## Abstract

In prior work on multiprocessor real-time locking protocols, only protocols within the RNLP family support unrestricted lock nesting while guaranteeing asymptotically optimal priority-inversion blocking bounds. However, these protocols support nesting at the expense of increasing the cost of processing non-nested lock requests, which tend to be the common case in practice. To remedy this situation, a new *fast-path mechanism* is presented herein that extends prior RNLP variants by ensuring that non-nested requests are processed efficiently. This mechanism yields overhead and blocking costs for such requests that are nearly identical to those seen in the most efficient single-resource locking protocols. In experiments, the proposed fast-path mechanism enabled observed blocking times for non-nested requests that were up to 18 times lower than under an existing RNLP variant and improved schedulability over that variant and a simple group lock.

## Additional Graphs

Here we include the results from our protocol testing and schedulability study corresponding to the paper of this title, which can be found online.<sup>1</sup>

To evaluate both variants of our new protocol, we explored the tradeoffs between overhead and blocking via user-space experiments under a variety of scenarios. In our experiments, we varied the number of tasks (n), nesting depth ( $\mathbb{D} = |D_i|$ ), critical-section length  $(L_i)$ , probability of a request being nested (rather than non-nested), and probability of a request being a read request (rather than a write request). Each task was pinned to a single core, and for task counts of up to 18, all tasks were assigned to the same socket. To simulate behavior that would generate the worst-case lock overhead and blocking times, each task was configured to issue lock and unlock calls 10,000 times, as fast as possible. Each such lock-unlock call pair corresponded to a single request that was randomly chosen to be nested (or non-nested) and a read (or a write) given the scenario's parameters. In our graphs, we plot these worst-case values, which were obtained by computing the 99<sup>th</sup> percentile of all recorded results in order to filter out any spurious measurements.

We also conducted a large-scale overhead-aware schedulability study to explore the impact of the worstcase acquisition delays on system schedulability. We used SchedCAT, an open-source real-time schedulability test toolkit, to randomly generate task systems, implement blocking bound computations, and check for schedulability on an 18-core platform with global EDF scheduling. We generated task sets with specified system utilizations, task periods lengths, and per-task utilizations. Tasks were chosen to issue a single request with a given probability, and the probability of a given request being a read (as opposed to a write) or nested (as opposed to non-nested) was specified. For each generated task system, we computed the impact of blocking on each request given the constraints discussed above. We did not consider most types of overhead, such as migration overhead, release blocking, and other overhead sources that impact each scheme similarly. However, we did account for the overhead incurred by using a specific locking protocol; we applied the appropriate overhead values for nested and non-nested read and write requests based on the prevoiusly described experiments. We chose the maximum values of lock and unlock overhead from relevant scenarios.

<sup>\*</sup>Work supported by NSF grants CNS 1409175, CPS 1446631, CNS 1563845, and CNS 1717589, AFOSR grant FA9550-14-1-0161, ARO grant W911NF-14-1-0499, ARO grant W911NF-17-1-0294, and funding from General Motors. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGS 1650116. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

<sup>&</sup>lt;sup>1</sup>See paper online: http://www.cs.unc.edu/~anderson/papers.html.



Figure 1: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the PF-TL and the fast RW-RNLP variants. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2.

Figure 2: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the PF-TL and the fast RW-RNLP variants. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8.



Figure 3: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the PF-TL and the fast RW-RNLP variants. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2.

Figure 4: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the PF-TL and the fast RW-RNLP variants. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8.



Figure 5: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the PF-TL and the fast RW-RNLP variants. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2.

Figure 6: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the PF-TL and the fast RW-RNLP variants. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8.



Figure 7: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the PF-TL and the fast RW-RNLP variants. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2.

Figure 8: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the PF-TL and the fast RW-RNLP variants. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8.



Figure 9: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the PF-TL and the fast RW-RNLP variants. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2.

Figure 10: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the PF-TL and the fast RW-RNLP variants. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8.



Figure 11: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the PF-TL and the fast RW-RNLP variants. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 20\mu$ s,  $n_r =$ 64, and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown.

Figure 12: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the PF-TL and the fast RW-RNLP variants. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown.



Figure 13: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the PF-TL and the fast RW-RNLP variants. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown.

Figure 14: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the PF-TL and the fast RW-RNLP variants. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown.



Figure 15: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the PF-TL and the fast RW-RNLP variants. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown.

Figure 16: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the PF-TL and the fast RW-RNLP variants. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown.





Figure 17: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.

Figure 18: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.



Figure 19: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 20: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 20\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 22: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 20\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(a) Lock overhead.



(c) Blocking

15 20 Number of Tasks 25

10

35





Figure 24: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu$ s,  $n_r =$ 64, and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 25: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.



Figure 26: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i =$  $20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.



Figure 27: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 28: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 20\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 29: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 30: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 20\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(a) Lock overhead.





Figure 31: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 32: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu$ s,  $n_r =$ 64, and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 34: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.







Figure 36: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 60\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 38: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 60\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(c) Blocking





Figure 40: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu$ s,  $n_r =$ 64, and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







20 Number of Tasks (a) Lock overhead. W-NN - Fast RW-RNLP (RW-RNLP R-NN - Fast RW-RNLP (RW-RNLP 5 20 Number of Tasks

(b) Unlock overhead.









Figure 44: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 60\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 45: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 46: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 60\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(a) Lock overhead.





Figure 47: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 48: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu$ s,  $n_r =$ 64, and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 49: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu$ s,  $n_\tau = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.



Figure 50: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.



Figure 51: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100 \mu s, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



25

to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 54: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 100\mu s, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(a) Lock overhead.









Figure 56: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu$ s,  $n_r =$ 64, and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 58: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i =$  $100\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.



Figure 59: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 60: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 100\mu s, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 62: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 100\mu s, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(a) Lock overhead.



100

10



W-N - Fast RW-RNLP (RW-RNLP\* R-N - Fast RW-RNLP (RW-RNLP\*)

5 20 Number of Tasks

(a) Lock overhead.

зс

35

35

30

25

15

W-N - Fast RW-RNLP (RW-RNLP\*) R-N - Fast RW-RNLP (RW-RNLP\*)

W-N - RW-RNLP

R-N - RW-RNLF

•

(spu

1.5

· ↔ · W-N - RW-RNLP • ■ · R-N - RW-RNLP

25

20

ĕ 15

Lock

Figure 63: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100 \mu s$ ,  $n_r = 64$ , and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

5 20 Number of Tasks

(c) Blocking.

Figure 64: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100 \mu s$ ,  $n_r =$ 64, and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



 W-NN - Fast RW-RNLP (RW-RNLP\*)
R-NN - Fast RW-RNLP (RW-RNLP\*)
W-NN - RW-RNLP
R-NN - RW-RNLP ଞ୍ଚ 50 4( Sp 30 8 20 Lock 10 20 Number of Tasks (a) Lock overhead. W-NN - Fast RW-RNLP (RW-RNLP • R-NN - Fast RW-RNLP (RW-RNLP 30 W-NN - RW-RNLP R-NN - RW-RNLF 25 20 eads 15 Unlock Ov 10 20 mber of Tasks (b) Unlock overhead. 200 W-NN - Fast RW-RNLP (RW-RNLP\* 4 . R-NN - Fast RW-RNLP (RW-RNLP W-NN - RW-RNLP 175 R-NN - RW-RNLE 150 (spi 125 100 Blocking ( 75 50 25 10 15 20 Number of Tasks 25 30 35 (c) Blocking.

W-NN - Fast RW-RNLP (RW-RNLP

Figure 65: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.

Figure 66: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i =$  $20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.







Figure 68: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 20\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 69: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 70: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 20\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.






Figure 72: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu$ s,  $n_r =$ 64, and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 74: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.



Figure 75: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 76: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 20\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 78: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 20\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 79: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 80: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu$ s,  $n_r =$ 64, and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 82: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.



Figure 83: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu s, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



W-N - Fast RW-RNLP (RW-RNLP\*) R-N - Fast RW-RNLP (RW-RNLP\*) W-NN - Fast RW-RNLP (RW-RNLP

R-NN - Fast RW-RNLP (RW-RNLP\*

30

5 20 Number of Tasks

(c) Blocking.

W-N - RW-RNLF

R-N - RW-RNLP

25

R-NN - RW-RNLF

-•

•

5 20 Number of Tasks

5 20 Number of Tasks



Figure 85: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 86: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 60\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





2000

Ē 1500

B 1000

50

10



W-N - Fast RW-RNLP (RW-RNLP\*)

12

Figure 87: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

5 20 Number of Tasks

(c) Blocking

Figure 88: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu$ s,  $n_r =$ 64, and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

35

30







Figure 90: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.







Figure 92: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 60\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 94: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 60\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 95: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 96: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu$ s,  $n_r =$ 64, and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 98: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i =$  $100\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.



Figure 99: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 100: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 100\mu s, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 101: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 102: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 100\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







llocking

1000

10



W-N - Fast RW-RNLP (RW-RNLP\* R-N - Fast RW-RNLP (RW-RNLP\*)

W-N - RW-RNLF W-N - RW-RNL

-0-3.5

10

W-N - Fast RW-RNLP (RW-RNLP\*) R-N - Fast RW-RNLP (RW-RNLP\*)

15

Number

(a) Lock overhead.

20 of Tasks

12

10

(spc

eads

Lock Over

(c) Blocking.

Figure 103: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100 \mu s$ ,  $n_r = 64$ , and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

5 20 Number of Tasks

(c) Blocking

Figure 104: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100 \mu s$ ,  $n_r =$ 64, and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

35

30



Figure 105: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu$ s,  $n_\tau = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.

Figure 106: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i =$  $100\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.





→ W-N - Fast RW-RNLP (RW-RNLP\* R-N - Fast RW-RNLP (RW-RNLP\*)

W-NN - Fast RW-RNLP (RW-RNLP\*

R-NN - Fast RW-RNLP (RW-RNLP\* W-N - RW-RNLP R-N - RW-RNLP R-NN - RW-RNLP

25

Ĕ 15

Lock Overhe

•

Figure 107: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(c) Blocking.





Figure 110: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 100\mu s, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 111: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 112: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu$ s,  $n_r =$ 64, and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 113: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.

Figure 114: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.



Figure 115: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 116: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 20\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 118: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 20\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





(c) Blocking.

5 20 Number of Tasks

30

10





Figure 120: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu s$ ,  $n_r =$ 64, and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 121: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu s$ ,  $n_\tau = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.



(c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.







Figure 124: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 20\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 125: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 126: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 20\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 127: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 128: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 20\mu$ s,  $n_r =$ 64, and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 129: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu$ s,  $n_\tau = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.



Figure 130: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.



Figure 131: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 132: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 60\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 133: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 134: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 60\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 135: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 136: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu$ s,  $n_r =$ 64, and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 138: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.



Figure 139: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 140: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 60\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 142: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 60\mu$ s,  $n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.






Figure 143: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 144: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 60\mu$ s,  $n_r =$ 64, and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 145: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100 \mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.



Figure 146: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i =$  $100\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.



Figure 147: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100 \mu s, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



- RW-RNLE



Figure 149: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 150: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 100\mu s, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 151: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 152: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu$ s,  $n_r =$ 64, and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 154: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.







Figure 156: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 100\mu s, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 158: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, L_i = 100\mu s, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.











Figure 160: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ ,  $L_i = 100\mu$ s,  $n_r =$ 64, and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 161: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r =$ 64, and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.

Figure 162: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.



Figure 163: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 164: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 165: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 166: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.









Figure 168: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ , and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 169: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r =$ 64, and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.

Figure 170: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.



Figure 171: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 172: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 173: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 174: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 175: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 18,  $n_{r} = 64$ , and  $|D_{i}| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 176: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ , and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 177: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r =$ 64, and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.

Figure 178: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.



Figure 179: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 180: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 181: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 182: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 184: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ , and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 185: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r =$ 64, and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.

Figure 186: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.



Figure 187: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 188: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 189: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 190: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 191: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 18,  $n_{r} = 64$ , and  $|D_{i}| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 192: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ , and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 193: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r =$ 64, and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.

Figure 194: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.



Figure 195: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 196: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 197: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 198: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.









Figure 200: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ , and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 201: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r =$ 64, and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.

Figure 202: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.



Figure 203: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 204: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 205: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 206: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 208: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $n_r = 64$ , and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 209: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.

Figure 210: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.



Figure 211: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 212: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 213: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 214: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.






Figure 216: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 217: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.

Figure 218: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.







Figure 220: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 221: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 222: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 223: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 36,  $n_{r} = 64$ , and  $|D_{i}| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 224: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 225: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.

Figure 226: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.



Figure 227: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 228: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 229: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 230: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 231: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 36,  $n_{r} = 64$ , and  $|D_{i}| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 232: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 233: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.

Figure 234: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.



Figure 235: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 236: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 237: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $m = 36, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 238: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 240: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 241: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.

Figure 242: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.



Figure 243: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 244: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 246: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 248: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(c) Blocking.



Figure 250: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.



Figure 251: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 252: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 253: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 254: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, n_r = 64, |D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 255: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 36,  $n_{r} = 64$ , and  $|D_{i}| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

출 음 2000

1000

는 원 3000

Figure 256: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $n_r = 64$ , and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

100

an

60

CS Lenath

(c) Blocking.



(c) Blocking.



Figure 257: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i =$  $20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.



(a) Lock overhead.



(b) Unlock overhead.



14 କ୍<u>ଥି</u> 12 W-N - Fast RW-RNLP (RW-RNLP\* 8 10 - R-N - Fast RW-RNLP (RW-RNLP) W-NN - Fast RW-RNLP (RW-RNL) • R-NN - Fast RW-RNLP (RW-RNLP\* W-N - RW-RNLP • R-N - RW-RNLP R-NN - RW-RNLP

Lock Overheads (mic

(a) Lock overhead.

. Percent Reads



(b) Unlock overhead.



Figure 259: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $m = 18, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 260: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 262: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



· ↔ · W-N - RW-RNLP • ■ · R-N - RW-RNLP 12 10 Lock Ove 60 80 100 cent Reads (a) Lock overhead. 2.5 W-N - Fast RW-RNLP (RW-RNLP\* R-N - Fast RW-RNLP (RW-RNLP\* W-N - RW-RNLP 20 ds) R-N - RW-RNLP 1.5 Unlock Overheads Percent Reads (b) Unlock overhead. 700 600 (spi 500 W-N - Fast RW-RNLP (RW-RNLP\* R-N - Fast RW-RNLP (RW-RNLP\* 400 W-N - RW-RNLP 300 R-N - RW-RNLI ĕ 200 100 20 60 80 100 40 € Percent Reads (c) Blocking.

W-N - Fast RW-RNLP (RW-RNLP\* R-N - Fast RW-RNLP (RW-RNLP\*

Figure 263: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 18,  $L_{i} = 20\mu$ s,  $n_{r} = 64$ , and  $|D_{i}| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 264: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(c) Blocking.



Figure 265: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i =$  $60\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.





(b) Unlock overhead.







(a) Lock overhead.



(b) Unlock overhead.



Figure 268: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 269: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 270: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 272: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(c) Blocking.



Figure 273: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i =$  $100\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.



(a) Lock overhead.



(b) Unlock overhead.





(a) Lock overhead.



(b) Unlock overhead.



Figure 275: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 276: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 100\mu \text{s}, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 278: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 100\mu \text{s}, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 280: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(c) Blocking.



Figure 281: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i =$  $20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.



(a) Lock overhead.



(b) Unlock overhead.



Figure 283: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 284: (a) Lock and (b) unlock overheads

and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(a) Lock overhead.



(b) Unlock overhead.





(a) Lock overhead.



(b) Unlock overhead.



Figure 285: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 286: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.




Figure 287: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 18,  $L_{i} = 20\mu$ s,  $n_{r} = 64$ , and  $|D_{i}| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 288: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(c) Blocking.



Figure 289: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i =$  $60\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.





(b) Unlock overhead.







Figure 292: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





(b) Unlock overhead.







Figure 294: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 296: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(c) Blocking.

with probability 0.



Figure 297: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i =$  $100\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request request with probability 0.





Figure 299: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

40 60 Percent Reads

(c) Blocking.

W-N - RW-RNLP

R-N - RW-RNLP

R-NN - RW-RNL

80

100

1500

500

20

훕 1000







(b) Unlock overhead.





14

•

0.75

R-NN - Fast RW-RNLP (RW-RNLP

W-N - RW-RNLF



Figure 301: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $m = 18, L_i = 100 \mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







14

Figure 303: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100 \mu s, n_r = 64, and |D_i| = 4.$  Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 304: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i =$  $100\mu s, n_r = 64$ , and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(c) Blocking.

with probability 0.



Figure 305: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i =$  $20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.



(a) Lock overhead.



(b) Unlock overhead.



Figure 307: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 308: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





(b) Unlock overhead.



Figure 309: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 310: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 311: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 18,  $L_{i} = 20\mu$ s,  $n_{r} = 64$ , and  $|D_{i}| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 312: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(c) Blocking.



Figure 313: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i =$  $60\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.



(a) Lock overhead.



(b) Unlock overhead.







Figure 316: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(a) Lock overhead.



(b) Unlock overhead.



Figure 317: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 318: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 319: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 18,  $L_{i} = 60\mu$ s,  $n_{r} = 64$ , and  $|D_{i}| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 320: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(c) Blocking.



Figure 321: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i =$  $100\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.





(b) Unlock overhead.







Figure 324: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 100\mu \text{s}, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(a) Lock overhead.



(b) Unlock overhead.







Figure 326: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 100\mu$ s,  $n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 327: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 328: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 329: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.

Figure 330: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.



(a) Lock overhead.







(a) Lock overhead.



(b) Unlock overhead.



Figure 331: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 332: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





(b) Unlock overhead.







Figure 333: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 334: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 336: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(c) Blocking.



Figure 337: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i =$  $60\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.





(b) Unlock overhead.









Figure 339: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 340: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(a) Lock overhead.



(b) Unlock overhead.







(b) Unlock overhead.



Figure 341: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 342: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





W-N - Fast RW-RNLP (RW-RNLP\* R-N - Fast RW-RNLP (RW-RNLP\*)

25

Figure 343: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 36,  $L_{i} = 60\mu$ s,  $n_{r} = 64$ , and  $|D_{i}| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





(c) Blocking.



Figure 346: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.

100

100



(a) Lock overhead.









(b) Unlock overhead.



Figure 347: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 348: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 100\mu \text{s}, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(a) Lock overhead.



(b) Unlock overhead.





(c) Blocking.

Figure 349: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 350: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 100\mu \text{s}, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 351: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 352: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i| = 2$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 353: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.

Figure 354: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.



(a) Lock overhead.











Figure 355: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 356: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(a) Lock overhead.







(a) Lock overhead.



1400 1200 W-N - Fast RW-RNLP (RW-RNLP\* 1000 - R-N - Fast RW-RNLP (RW-RNLP\* W-NN - Fast RW-RNLP (RW-RNLP 800 R-NN - Fast RW-RNLP (RW-RNLP • W-N - RW-RNLF Blocking ( 600 . R-N - RW-RNLP R-NN - RW-RNLI 400 200 60 10 6 Percent Reads (c) Blocking.

Figure 357: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 358: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.






Figure 360: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(c) Blocking.



Figure 361: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i =$  $60\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.











Figure 364: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(a) Lock overhead.







(a) Lock overhead.



(b) Unlock overhead.



Figure 365: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.









Figure 368: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(c) Blocking.



Figure 369: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i =$  $100\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.

Figure 370: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.



(a) Lock overhead.







(a) Lock overhead.



(b) Unlock overhead.



Figure 371: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 372: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 100\mu \text{s}, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(a) Lock overhead.







(a) Lock overhead.



(b) Unlock overhead.



Figure 373: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 374: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 100\mu \text{s}, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 375: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 376: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i| = 4$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 377: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.

Figure 378: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu s$ ,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.



(a) Lock overhead.







(c) Blocking.

Figure 379: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 380: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(a) Lock overhead.









Figure 381: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 382: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 383: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 36,  $L_{i} = 20\mu$ s,  $n_{r} = 64$ , and  $|D_{i}| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 384: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(c) Blocking.



Figure 386: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu s, n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.











Figure 387: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 388: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(a) Lock overhead.









Figure 389: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 390: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 391: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 36,  $L_{i} = 60\mu$ s,  $n_{r} = 64$ , and  $|D_{i}| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 392: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(c) Blocking.



Figure 393: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.

Figure 394: (a) Lock and (b) unlock overheads and (c) blocking for non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i| = 1$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.



(a) Lock overhead.







Figure 395: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 396: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 100\mu \text{s}, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

(c) Blocking.

40 60 Percent Reads



(a) Lock overhead.







Figure 397: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.







Figure 399: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 400: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i| = 6$ . Each request was randomly chosen to be a read (as opposed to a write) with probability as shown and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





(b) Unlock overhead.



Figure 401: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.



Figure 402: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.





60

80

100





Figure 404: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.







Figure 405: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.



Figure 406: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.













Figure 408: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.





(b) Unlock overhead.



Figure 409: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.



Figure 410: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 100\mu \text{s}, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.





Lock Overheads (micro





Figure 411: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.

40 60 Percent Nested

(c) Blocking.

100

500









(b) Unlock overhead.







Figure 414: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.





(b) Unlock overhead.



Figure 415: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.



Figure 416: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.







(b) Unlock overhead.







Figure 418: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.







(b) Unlock overhead.



Figure 419: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.



Figure 420: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.













Figure 422: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 100\mu \text{s}, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.











Figure 424: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 100\mu \text{s}, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.









Figure 425: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.



Figure 426: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.







(b) Unlock overhead.



Figure 427: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.



Figure 428: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.









Figure 429: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.



Figure 430: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.






(b) Unlock overhead.





Figure 431: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.





(a) Lock overhead.



(b) Unlock overhead.



Figure 433: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.



Figure 434: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 100\mu \text{s}, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.







(b) Unlock overhead.







(b) Unlock overhead.



Figure 435: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.







Figure 437: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.



Figure 438: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.



(a) Lock overhead.









Figure 440: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.







Figure 441: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.















Figure 444: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.







Figure 446: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 100\mu$ s,  $n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.













Figure 448: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 100\mu \text{s}, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 2$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.











Figure 449: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.

Figure 450: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.





(c) Blocking.

Figure 451: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.





(a) Lock overhead.



(b) Unlock overhead.





(c) Blocking.

Figure 453: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.

Figure 454: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.







Figure 456: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.







(b) Unlock overhead.







(b) Unlock overhead.



Figure 457: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.







(c) Blocking.

Figure 459: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 4$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.















(b) Unlock overhead.



Figure 461: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.

Figure 462: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.





(a) Lock overhead.



(b) Unlock overhead.



Figure 463: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.

Figure 464: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 20\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.











Figure 465: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.

Figure 466: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 60\mu \text{s}, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.









(b) Unlock overhead.



Figure 467: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for non-nested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.

Figure 468: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 60\mu s, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.











(a) Lock overhead.



(b) Unlock overhead.



Figure 469: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.

Figure 470: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 100\mu \text{s}, n_r = 64, |D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.



(b) Unlock overhead.



(a) Lock overhead.



(b) Unlock overhead.



Figure 471: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ ,  $|D_i| = 1$  for nonnested requests, and  $|D_i| = 6$  for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability as shown. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP if some requests were nested, as read requests can access any resource.





(c) Blocking. Figure 473: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ ,  $m = 18, L_i = 20\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 474: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 20\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 475: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 476: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 20\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 477: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 18,  $L_{i} = 20\mu$ s,  $n_{r} = 64$ , and  $|D_{i}|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 478: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



12 (spuo: 10 W-N - Fast RW-RNLP (RW-RNLP\* R-N - Fast RW-RNLP (RW-RNLP\*) W-NN - Fast RW-RNLP (RW-RNLP Overheads (micro 8 R-NN - Fast RW-RNLP (RW-RNLP\* 6 W-N - RW-RNLP R-N - RW-RNLP • R-NN - RW-RNLF Lock 0 5 6 Level of Ne (a) Lock overhead. 1.6 econds) 1.2 W-N - Fast RW-RNLP (RW-RNLP\*) - R-N - Fast RW-RNLP (RW-RNLP\* ມັ 1.0 W-NN - Fast RW-RNLP (RW-RNLF 0.0 overheads ( W-N - RW-RNLP R-N - RW-RNLP -R-NN - RW-RNI F yooln 0.4 0.2 (b) Unlock overhead. 700 600 500 W-N - Fast RW-RNLP (RW-RNLP\*) - R-N - East RW-RNI P (RW-RNI P\* W-NN - Fast RW-RNLP (RW-RNLP 400 4 • \*\*\* R-NN - Fast RW-RNLP (RW-RNLP W-N - RW-RNLF 문 300 및 R-N - RW-RNLP Block R-NN - RW-RNLE 200 100 Level of Nest (c) Blocking.

14

Figure 479: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 480: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 20\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 481: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 482: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 20\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 483: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 18,  $L_{i} = 20\mu$ s,  $n_{r} = 64$ , and  $|D_{i}|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 484: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP<sup>\*</sup>. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 485: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 486: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 60\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 487: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 488: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 60\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Lock Overheads (microseconds) W-N - Fast RW-RNLP (RW-RNLP R-N - Fast RW-RNLP (RW-RNLP\* W-N - RW-RNLP R-N - RW-RNLP R-N - RW-RNLP 5 6 Level of Nestedne (a) Lock overhead. 2.00 1.75 (SP 1.50 1.25 W-N - Fast RW-RNLP (RW-RNLP\* W-N - Fast RW-RNLP (*RW-RNLP* R-N - Fast RW-RNLP (*RW-RNLP* W-N - RW-RNLP R-N - RW-RNLP • ) 1.00 Overheads 0.75 \* 0.50 0.25 5 6 Level of Nestedn (b) Unlock overhead. 175 1500 ឆ្នាំ 1250 W-N - Fast RW-RNLP (RW-RNLP\* 1000 R-N - Fast RW-RNLP (RW-W-N - RW-RNLP R-N - RW-RNLP 750 Block 500 250 (c) Blocking.

Figure 489: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 18,  $L_{i} = 60\mu$ s,  $n_{r} = 64$ , and  $|D_{i}|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 490: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 491: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 492: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 60\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 493: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 494: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 60\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 495: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 18,  $L_{i} = 60\mu$ s,  $n_{r} = 64$ , and  $|D_{i}|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 496: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 497: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu s$ ,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 498: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 100\mu \text{s}, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 499: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu s$ ,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 500: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 501: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 502: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP<sup>\*</sup>. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.


14 12 10 W-N - Fast RW-RNLP (RW-RNLP R-N - Fast RW-RNLP (RW-RNLP\*) W-NN - Fast RW-RNLP (RW-RNLP Lock Overheads (mic R-NN - Fast RW-RNLP (RW-RNLP\* W-N - RW-RNLP R-N - RW-RNLP e . R-NN - RW-RNLF 5 6 Level of Ne (a) Lock overhead. 1.6 දු 1.4 ğ 1.2 W-N - Fast RW-RNLP (RW-RNLP\* - R-N - Fast RW-RNLP (RW-RNLP\* Ĕ 1.0 W-NN - Fast RW-RNLP (*RW-RNLP* R-NN - Fast RW-RNLP (*RW-RNLP* W-N - RW-RNLP
R-N - RW-RNLP
C.NN - RW-RNLP W-NN - Fast RW-RNLP (RW-RNL) 8.0 overheads ( 0.0 Unlock 0.2 Level of Nest (b) Unlock overhead. 350 300 <u>ශ</u> 2500 W-N - Fast RW-RNLP (RW-RNLP\*) R-N - East RW-RNI P (RW-RNI P\* 2000 W-NN - Fast RW-RNLP (RW-RNLP • R-NN - Fast RW-RNLP (RW-RNLP 1500 W-N - RW-RNLF R-N - RW-RNLF 8 100C R-NN - RW-RNLE 500 Level of Nest (c) Blocking.

Figure 503: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 504: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 100\mu \text{s}, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 505: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu s$ ,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 506: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 18, L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 507: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 508: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 18,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 509: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 510: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 20\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 511: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 512: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 20\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 513: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 36,  $L_{i} = 20\mu$ s,  $n_{r} = 64$ , and  $|D_{i}|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 514: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP<sup>\*</sup>. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 515: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 516: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 20\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 517: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 518: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 20\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



(Sp 3.5 3.0 W-N - Fast RW-RNLP (RW-RNLP\*) 0 overheads (mio R-N - Fast RW-RNLP (RW-RNLP • W-N - RW-RNLP R-N - RW-RNLF 90 1.5 1.0 0.5 10 Level of Nest (b) Unlock overhead. 140 1200 1000 W-N - Fast RW-RNLP (RW-RNLP\* R-N - Fast RW-RNLP (RW-RNL) W-N - RW-RNLP R-N - RW-RNLP 800 0 600 Blocking 400 200 10 Level of Nester (c) Blocking. Figure 520: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 20\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to

write expansion,  $|D_i|$  was inflated to 64 for all write

requests under the RW-RNLP, as read requests can

W-N - Fast RW-RNLP (RW-RNIP

R-N - Fast RW-RNLP (RW-RNLP\* W-N - RW-RNLP

R-N - RW-RNLI

5 6 Level of Nestednes

(a) Lock overhead.

25

20

Overheads (r

Lock

4.0

Figure 519: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 36,  $L_{i} = 20\mu$ s,  $n_{r} = 64$ , and  $|D_{i}|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

access any resource.



Figure 521: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 522: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 60\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 523: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 524: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 60\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



12 10 W-N - Fast RW-RNLP (RW-RNLP\* R-N - Fast RW-RNLP (RW-RNLP\*) W-N - RW-RNLP R-N - RW-RNLP • Lock Overheads ( 5 6 Level of Nestedne (a) Lock overhead. W-N - Fast RW-RNLP (RW-RNLP\*) R-N - Fast RW-RNLP (RW-RNLP\*) • W-N - RW-RNLP Coverheads (microseconds) R-N - RW-RNLF Unlock (b) Unlock overhead. 350 300 ភ្នំ 2500 W-N - Fast RW-RNLP (RW-RNLP\* 2000 R-N - Fast RW-RNLP (RW-W-N - RW-RNLP R-N - RW-RNLP R-N - RW-RNLP 1500 Bloc 1000 500 10 Level of Nest (c) Blocking.

Figure 525: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 36,  $L_{i} = 60\mu$ s,  $n_{r} = 64$ , and  $|D_{i}|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 526: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP<sup>\*</sup>. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 527: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 528: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 60\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 529: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 530: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 60\mu s, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 531: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the  $R^{3}LP$ . Here, for each request  $\mathcal{R}_{i}$ , m = 36,  $L_{i} = 60\mu$ s,  $n_{r} = 64$ , and  $|D_{i}|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_{i}|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 532: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP<sup>\*</sup>. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 60\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



·····×····\*

Figure 533: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu s$ ,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 534: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 100\mu \text{s}, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 535: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu s$ ,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 536: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 100\mu \text{s}, n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



12 (spug 10 W-N - Fast RW-RNLP (RW-RNLP\* www-rast RW-RNLP (*RW-RNLP*\*) R-N - Fast RW-RNLP (*RW-RNLP*\*) W-N - RW-RNLP R-N - RW-RNLP • Lock Overheads ( 10 5 6 Level of Nestedne (a) Lock overhead. W-N - Fast RW-RNLP (RW-RNLP\*) R-N - Fast RW-RNLP (RW-RNLP\*) • W-N - RW-RNLP Unlock Overheads (microseconds) N W A V R-N - RW-RNLF (b) Unlock overhead. 5000 W-N - Fast RW-RNLP (RW-RNLP\* R-N - Fast RW-RNLP (RW-RNL) W-N - RW-RNLP R-N - RW-RNLP 3000 Slockin 200 1000 10 Level of Neste (c) Blocking.

Figure 537: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 538: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.2 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 539: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu s$ ,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 540: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.2. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 541: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu s$ ,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 542: (a) Lock and (b) unlock overheads and (c) blocking for nested and non-nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP\*. Here, for each request  $\mathcal{R}_i, m = 36, L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 0.8. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.





Figure 543: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the R<sup>3</sup>LP. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.

Figure 544: (a) Lock and (b) unlock overheads and (c) blocking for nested read and write requests under the RW-RNLP and the fast RW-RNLP with the RW-RNLP<sup>\*</sup>. Here, for each request  $\mathcal{R}_i$ , m = 36,  $L_i = 100\mu$ s,  $n_r = 64$ , and  $|D_i|$  is as shown for nested requests. Each request was randomly chosen to be a read (as opposed to a write) with probability 0.8 and to be a nested request with probability 1. Due to write expansion,  $|D_i|$  was inflated to 64 for all write requests under the RW-RNLP, as read requests can access any resource.



Figure 545: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, short critical-section lengths, and read probability 0 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 546: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, short critical-section lengths, and read probability 0 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 547: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, short critical-section lengths, and read probability 0 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 548: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, short critical-section lengths, and read probability 0 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 549: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 550: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 551: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 552: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, short critical-section lengths, and read probability 0.2 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 553: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 554: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 555: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 556: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, short critical-section lengths, and read probability 0.5 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 557: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 558: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 559: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.


Figure 560: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, short critical-section lengths, and read probability 0.8 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 561: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, short critical-section lengths, and read probability 0 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 562: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, short critical-section lengths, and read probability 0 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 563: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, short critical-section lengths, and read probability 0 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 564: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, short critical-section lengths, and read probability 0 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 565: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 566: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 567: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 568: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, short critical-section lengths, and read probability 0.2 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 569: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 570: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 571: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 572: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, short critical-section lengths, and read probability 0.5 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 573: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 574: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 575: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 576: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, short critical-section lengths, and read probability 0.8 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 577: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, short critical-section lengths, and read probability 0 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 578: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, short critical-section lengths, and read probability 0 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 579: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, short critical-section lengths, and read probability 0 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 580: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, short critical-section lengths, and read probability 0 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 581: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 582: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 583: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 584: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, short critical-section lengths, and read probability 0.2 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 585: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 586: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 587: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 588: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, short critical-section lengths, and read probability 0.5 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 589: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 590: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 591: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 592: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, short critical-section lengths, and read probability 0.8 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 593: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, long critical-section lengths, and read probability 0 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 594: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, long critical-section lengths, and read probability 0 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 595: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, long critical-section lengths, and read probability 0 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.


Figure 596: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, long critical-section lengths, and read probability 0 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 597: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 598: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 599: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 600: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, long critical-section lengths, and read probability 0.2 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 601: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 602: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 603: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 604: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, long critical-section lengths, and read probability 0.5 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 605: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 606: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 607: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 608: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, short periods, long critical-section lengths, and read probability 0.8 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 609: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, long critical-section lengths, and read probability 0 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 610: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, long critical-section lengths, and read probability 0 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 611: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, long critical-section lengths, and read probability 0 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 612: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, long critical-section lengths, and read probability 0 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 613: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 614: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 615: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 616: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, long critical-section lengths, and read probability 0.2 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 617: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 618: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 619: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 620: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, long critical-section lengths, and read probability 0.5 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 621: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 622: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 623: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 624: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, moderate periods, long critical-section lengths, and read probability 0.8 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 625: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, long critical-section lengths, and read probability 0 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 626: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, long critical-section lengths, and read probability 0 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 627: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, long critical-section lengths, and read probability 0 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 628: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, long critical-section lengths, and read probability 0 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 629: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 630: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 631: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.


Figure 632: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, long critical-section lengths, and read probability 0.2 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 633: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 634: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 635: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 636: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, long critical-section lengths, and read probability 0.5 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 637: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 638: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 639: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 640: Hard real-time schedulability results with varying nested probabilities for the scenario with medium task utilizations, long periods, long critical-section lengths, and read probability 0.8 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 641: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, short critical-section lengths, and read probability 0 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 642: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, short critical-section lengths, and read probability 0 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 643: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, short critical-section lengths, and read probability 0 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 644: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, short critical-section lengths, and read probability 0 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 645: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 646: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 647: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 648: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, short critical-section lengths, and read probability 0.2 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 649: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 650: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 651: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 652: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, short critical-section lengths, and read probability 0.5 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 653: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 654: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 655: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 656: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, short critical-section lengths, and read probability 0.8 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 657: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, short critical-section lengths, and read probability 0 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 658: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, short critical-section lengths, and read probability 0 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 659: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, short critical-section lengths, and read probability 0 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 660: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, short critical-section lengths, and read probability 0 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 661: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 662: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 663: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 664: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, short critical-section lengths, and read probability 0.2 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 665: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 666: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 667: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.


Figure 668: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, short critical-section lengths, and read probability 0.5 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 669: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 670: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 671: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 672: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, short critical-section lengths, and read probability 0.8 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 673: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, short critical-section lengths, and read probability 0 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 674: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, short critical-section lengths, and read probability 0 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 675: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, short critical-section lengths, and read probability 0 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 676: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, short critical-section lengths, and read probability 0 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 677: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 678: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 679: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, short critical-section lengths, and read probability 0.2 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 680: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, short critical-section lengths, and read probability 0.2 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 681: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 682: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 683: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, short critical-section lengths, and read probability 0.5 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 684: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, short critical-section lengths, and read probability 0.5 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 685: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 686: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 687: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, short critical-section lengths, and read probability 0.8 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 688: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, short critical-section lengths, and read probability 0.8 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 689: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, long critical-section lengths, and read probability 0 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 690: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, long critical-section lengths, and read probability 0 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 691: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, long critical-section lengths, and read probability 0 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 692: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, long critical-section lengths, and read probability 0 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 693: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 694: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 695: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 696: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, long critical-section lengths, and read probability 0.2 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 697: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 698: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 699: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 700: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, long critical-section lengths, and read probability 0.5 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 701: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 702: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 703: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.


Figure 704: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, short periods, long critical-section lengths, and read probability 0.8 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 705: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, long critical-section lengths, and read probability 0 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 706: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, long critical-section lengths, and read probability 0 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 707: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, long critical-section lengths, and read probability 0 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 708: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, long critical-section lengths, and read probability 0 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 709: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 710: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 711: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 712: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, long critical-section lengths, and read probability 0.2 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 713: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 714: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 715: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 716: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, long critical-section lengths, and read probability 0.5 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 717: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 718: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 719: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 720: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, moderate periods, long critical-section lengths, and read probability 0.8 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 721: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, long critical-section lengths, and read probability 0 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 722: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, long critical-section lengths, and read probability 0 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 723: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, long critical-section lengths, and read probability 0 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 724: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, long critical-section lengths, and read probability 0 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 725: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 726: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 727: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, long critical-section lengths, and read probability 0.2 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 728: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, long critical-section lengths, and read probability 0.2 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 729: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 730: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 731: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, long critical-section lengths, and read probability 0.5 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 732: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, long critical-section lengths, and read probability 0.5 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 733: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.1 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 734: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.2 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 735: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, long critical-section lengths, and read probability 0.8 in which a task has probability 0.5 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.



Figure 736: Hard real-time schedulability results with varying nested probabilities for the scenario with heavy task utilizations, long periods, long critical-section lengths, and read probability 0.8 in which a task has probability 1.0 of issuing a request. Nested probabilities are (a) 0.01, (b) 0.05, (c) 0.1, (d) 0.2, and (e) 0.5.