

# Asymmetry-aware Scalable Locking

Nian Liu<sup>†‡</sup>, Jinyu Gu<sup>†‡</sup>, Dahai Tang<sup>\*</sup>, Kenli Li<sup>\*</sup>, Binyu Zang<sup>†‡</sup>, Haibo Chen<sup>†‡</sup>

{nianliu, gujinyu}@sjtu.edu.cn, {seatang, lkl}@hnu.edu.cn, {byzang, haibo.chen}@sjtu.edu.cn

<sup>†</sup> Engineering Research Center for Domain-specific Operating Systems, Ministry of Education, China

<sup>‡</sup> Institute of Parallel and Distributed Systems (IPADS), Shanghai Jiao Tong University

<sup>\*</sup> College of Information Science and Engineering, Hunan University

## Abstract

The pursuit of power-efficiency is popularizing asymmetric multicore processors (AMP) such as ARM big.LITTLE, Apple M1 and recent Intel Alder Lake with big and little cores. However, we find that existing scalable locks fail to scale on AMP and cause collapses in either throughput or latency, or both, because their implicit assumption of symmetric cores no longer holds. To address this issue, we propose the first asymmetry-aware scalable lock named LibASL. LibASL provides a new lock ordering guided by applications' latency requirements, which allows big cores to reorder with little cores for higher throughput under the condition of preserving applications' latency requirements. Using LibASL only requires linking the applications with it and, if latency-critical, inserting few lines of code to annotate the coarse-grained latency requirement. We evaluate LibASL in various benchmarks including five popular databases on Apple M1. Evaluation results show that LibASL can improve the throughput by up to 5 times while precisely preserving the tail latency designated by applications.

**CCS Concepts:** • Software and its engineering → Mutual exclusion; Software performance; Multithreading.

**Keywords:** synchronization primitives, lock, scalability, asymmetric multicore processor

## 1 Introduction

Single-ISA asymmetric multicore processor (AMP) combines cores of different computing capacities in one processor [53, 54] and has been widely used in mobile devices (e.g., ARM big.LITTLE [19]). Combining both faster big cores and slower little cores together, AMP is more flexible in accommodating both performance-oriented and energy-efficiency-oriented scenarios, such as leveraging all cores to achieve

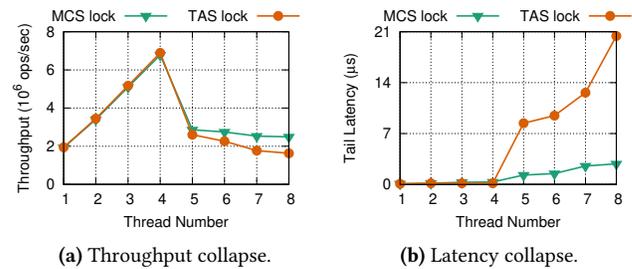
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

PPoPP '22, April 2–6, 2022, Seoul, Republic of Korea

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9204-4/22/04...\$15.00

<https://doi.org/10.1145/3503221.3508420>



**Figure 1.** Existing locks cause performance collapses on the Apple M1. M1 has 4 big and 4 little cores. The first 4 threads are bound to different big cores. Others are bound to different little cores. Threads are acquiring the same lock to read-modify-write 4 shared cache lines, and they will execute a fixed number ( $400 \times 2^2$ ) of NOP instructions between two lock acquisitions. The throughput is the number of executed critical sections in 1 second. The latency is the P99 tail latency from acquiring to releasing.

peak performance and using little cores only when energy is preferred. There is also a recent trend to embrace such an architecture in more general CPU processors, including the desktop and the edge server [2, 3, 12, 18]. As before, applications on AMP need to use locks for acquiring exclusive access to shared data. However, we observe that existing locks, including those scalable in the symmetric multicore processor (SMP) [26, 35] or non-uniform memory access system (NUMA) [28, 36–38, 50, 51, 59, 63], fail to scale in AMP and cause collapses in either throughput or latency, or both.

After an in-depth analysis, we find the main reason is that those locks (implicitly) assume symmetric cores, which does not hold on AMP. On the one side, locks that preserve lock acquisition fairness (i.e., give all cores an equal chance to lock), either short-term (e.g., MCS lock [61] passes the lock in a FIFO order) or long-term (e.g., NUMA-aware locks [28, 36–38, 50, 51, 59, 63] ensure the equal chance in a period), assume symmetric computing capacity. Therefore, in AMP, they give the slower little cores the same chance as the big cores to hold the lock, which introduces the longer execution time of the critical section in little cores to the critical path and causes throughput collapse. On the other side, locks that do not preserve the acquisition fairness rely on atomic operations to decide the lock holder (e.g., test-and-set spinlock). They assume a symmetric success rate of the atomic operation when executing simultaneously, which is also asymmetric in AMP. Thus, those locks are likely to be passed only among one type of core (i.e., either big cores or little cores), which

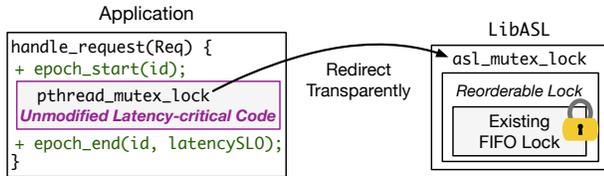


Figure 2. LibASL overview.

causes latency collapse even starvation to the others. Moreover, when the slower little cores have a higher chance to lock, the throughput also collapses due to the longer execution time of the critical sections on them. Figure 1 shows the performance collapses on Apple M1. Both the fair MCS lock and the unfair TAS (test-and-set) lock face throughput collapse when scaling to little cores. Besides, the TAS lock’s latency also collapses and is 6.2x longer than the MCS lock.

Facing the asymmetry in AMP, it is non-trivial to decide the lock ordering (who can lock firstly) for both high throughput and low latency. Binding threads only to big cores is an intuitive choice. However, using little cores can achieve higher throughput under a lower contention. Besides, it may violate the energy target as the energy-aware scheduler [9] could schedule threads to little cores for saving energy. Another intuitive approach is to give big cores a fixed higher chance to lock. However, the throughput and the latency are mutually exclusive in AMP. It is hard to find a one-size-fits-all (static) setting that can meet the application’s latency requirement and improve the throughput simultaneously.

In this paper, we propose an asymmetry-aware lock named LibASL as shown in Figure 2. Rather than ensuring the lock acquisition fairness that causes collapses on AMP, LibASL provides a new (dynamic) lock ordering guided directly by applications’ latency requirements to achieve better throughput under latency constraints. Atop of a FIFO waiting queue, LibASL allows big cores to reorder (lock before) with little cores as much as possible for higher throughput under the condition that the (reordered) victim will not miss the application’s latency target. To achieve such an ordering, we first design a *reorderable lock*, which exposes the reorder capability as a configurable time window. Big cores can only reorder with little cores during that time window. Atop of the *reorderable lock*, LibASL automatically chooses a suitable fine-grained reorder window according to applications’ coarse-grained latency requirements through a feedback mechanism. LibASL provides intuitive interfaces for developers to specify the coarse-grained latency requirements (e.g., request handling procedure) in the form of latency SLO (service level objective, e.g., 99% request should be complete within 50ms), which is widely adopted by both academia [46, 57, 68, 72] and industry [10, 33, 34]. To use LibASL, annotating the SLO is the **only** required effort if latency-critical, which is already clearly defined by applications in most cases. Non-latency-critical applications can benefit from LibASL without modifications.

We evaluate LibASL in multiple benchmarks including five popular databases on Apple M1, the only off-the-shelf desktop AMP yet. Results show that LibASL improves the throughput of `pthread_mutex_lock` by up to 5x (3.8x to MCS lock, 2.5x to TAS spinlock) while precisely maintaining the tail latency even in highly variable workloads.

In summary, this paper makes the following contributions:

- The first in-depth analysis of the performance collapses of existing locks on AMP.
- An asymmetry-aware scalable lock LibASL, which provides a new latency-SLO-guided lock ordering to achieve the best throughput the SLO allows on AMP.
- A thorough evaluation on the real desktop AMP (Apple M1) and real-world applications that confirms the effectiveness of LibASL.

## 2 Scalable Locking is Non-scalable on AMP

### 2.1 Asymmetric Multicore Processor

In this section, we introduce the major features of AMP.

First, the asymmetry in AMP is inherent. Although performance can also be asymmetric in SMP when using DVFS (dynamic voltage and frequency scaling), they can boost the frequency of the lagging core [24, 27, 67] while AMP cannot.

Second, asymmetric cores in recent AMP [4, 17, 18] are placed in one single cluster and share the same Last Level Cache (LLC). Thus, communication among cores in AMP is similar to SMP rather than NUMA.

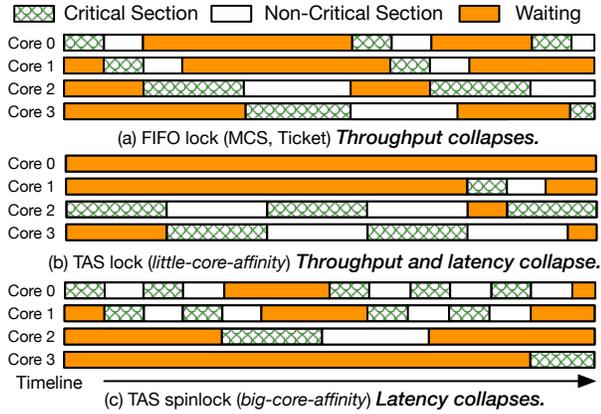
Third, the scheduler (e.g., the energy-aware scheduler in Linux [9]) can place different threads across asymmetric cores [40, 48]. Multi-threaded applications achieve better performance by leveraging all asymmetric cores [7] and need to use lock for synchronization among them as before.

### 2.2 A Study of Existing Scalable Locks in AMP

Scalable locking in SMP and NUMA has been extensively studied. However, existing scalable locks are non-scalable on AMP and encounter performance collapses. The main reason is that existing locks (implicitly) assume symmetric cores, which does not hold in AMP. There are two major differences between AMP and SMP that cause the collapses.

*First, the computing capacity is asymmetric.* Little cores spend a longer time executing the same critical section. Thus, when the lock preserves the acquisition fairness, it gives little cores an equal chance to hold the lock, which introduces the longer execution time of critical sections on them to the critical path and causes a throughput collapse.

We explain this problem through an example in Figure 3(a). The system includes two big cores (core 0/1) and two little cores (core 2/3), and they are competing for the same lock intensively. We divide the execution into three parts: executing the critical section, the non-critical section, and waiting for the lock. As shown in Figure 3(a), when ensuring the short-term (i.e., FIFO) lock acquisition fairness (e.g., the MCS



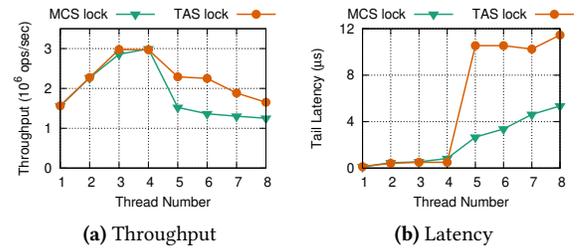
**Figure 3.** An example timeline (from left to right). Core 0/1 are big cores; 2/3 are little cores. More critical sections are executed means higher throughput; longer waiting time leads to longer latency.

lock [61]), threads hand over the lock in a FIFO order. As a result, the longer execution time of the critical section in core 2/3 will dominate the critical path and hurt the throughput.

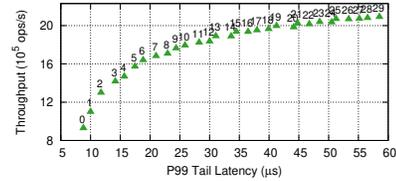
Besides the *short-term* acquisition fairness, previous work provides *long-term* acquisition fairness to improve the throughput in many-core processors or NUMA while keeping a relatively low latency, which also hurts the throughput in AMP. Malthusian lock [35] reduces the contention in the many-core processor for better throughput by blocking all competitors in the waiting queue except the head and the tail. It achieves long-term fairness by periodically shifting threads between blocking and acquiring. NUMA-aware locks [28, 36–38, 50, 51, 59, 63] batch the competitors from the same NUMA node to reduce the cross-node memory references. They achieve long-term fairness by periodically allowing different nodes to lock. All of these locks cannot scale in AMP. When splitting the asymmetric cores in AMP onto two different nodes, the long-term fairness will give the *little core nodes* an **equal** chance to lock as the *big core nodes*. Thus, similar to MCS, the longer execution time of critical sections in little cores dominate the critical path and causes throughput collapses.

**Implication 1:** *Lock ordering that respects acquisition fairness, either short-term or long-term, is no longer suitable in AMP.* In SMP or NUMA, preserving acquisition fairness can prevent starvation and achieve relatively low latency without degrading the throughput but causes collapses in AMP. Thus, a new lock ordering should be proposed for AMP to meet the latency goal while bringing higher throughput.

*Second, the success rate of atomic operations (e.g., test-and-set, TAS) is asymmetric.* On some AMP systems (e.g., ARM Kirin970 [11] and Intel L16G7 [13]), we observe that big cores have a stable advantage over little cores in winning the atomic TAS. While on other platforms (e.g., Apple M1),



**Figure 4.** When TAS lock shows big-core-affinity, it can achieve higher throughput but the latency still collapse.



**Figure 5.** Performance when setting different proportions (label on each point). N means big cores have N times higher chance to lock. The implementation of such an approach is explained in Section 4.

the advantage shifts between asymmetric cores<sup>1</sup>. Thus, locks that do not preserve acquisition fairness and rely on the atomic operation to decide the lock holder (e.g., TAS lock) also have the scalability issue in AMP. Such locks are likely to be held only by one type of core (i.e., either big cores or little cores), which causes a latency collapse even starvation to the others. Moreover, when the little cores have a bigger chance to hold the lock, the throughput also collapses due to the longer execution time of the critical sections on them.

As shown in Figure 3(b), when little cores show an advantage in winning the atomic TAS, they have more chance to lock (we name it as *little-core-affinity*). Thus, big cores can barely lock. Besides, the longer execution time of the critical sections on little cores will dominate the critical path and hurt the throughput. Similarly, when big cores show an advantage (*big-core-affinity*, Figure 3(c)), little cores will starve. Nevertheless, it allows big cores to lock before (reorder with) earlier little cores. Thus more critical sections are executed on faster big cores, which brings higher throughput.

We validate our observations on Apple M1. In Figure 1, we present the case when the TAS lock shows little-core-affinity. In Figure 4, we present another scenario when the TAS lock shows big-core-affinity<sup>2</sup>. In both scenarios, the fair MCS lock faces throughput collapses (over 50% degradation from 4 big cores to all cores), while the unfair TAS lock faces latency collapses. When the TAS lock shows little-core-affinity in Figure 1, its throughput also collapses and is 35% worse than

<sup>1</sup>On Apple M1, when executing TAS back-to-back (higher contention), little cores show a stable advantage. With the distance between two TAS increased (lower contention), big cores show a stable advantage.

<sup>2</sup>Specifically, read-modify-write 64 cache lines (4 cache lines in Figure 1) in the critical section. We identify the affinity by comparing the number of executed critical sections in different types of cores.

```

void request_handler(void) {
+ epoch_start(5); // Epoch 5 starts
  if (flag == PATH_1) {
    lock(&lock_1); // Critical Section 1
    lock & unlock (&lock_2); // Critical Section 2
    unlock(&lock_1);
  } else {
    lock & unlock (&lock_1); // Critical Section 3
  } // Unmodified Latency-critical Code
+ epoch_end(5, 1000); // Epoch 5 ends. SLO 1µs
}
    
```

Figure 6. Example of using LibASL.

the MCS lock when using all the cores. However, when the TAS lock shows big-core-affinity in Figure 4, more critical sections will be executed on the faster big cores, bringing 32% higher throughput than the MCS lock. The latency of the MCS lock also increases when scaling to little cores due to the longer execution time of critical sections on little cores. However, it is much shorter than the TAS lock. These observations still hold in real-world applications. When the TAS lock shows little-core-affinity in SQLite (detailed in Section 4.2), it has 49% worse throughput and 1.8x longer tail latency than the MCS lock. However, when the TAS lock shows big-core-affinity in UpscaleDB, it has 90% better throughput yet 2.5x longer tail latency than the MCS lock.

**Implication 2:** *Reordering to prioritize faster cores is indispensable in AMP for higher throughput, but it must be bounded.* When the TAS lock shows big-core-affinity, it reorders big cores with little cores unlimitedly and achieves higher throughput. However, the unbounded reordering causes a latency collapse. Thus, the reordering must be bounded for preserving applications’ latency requirements.

### 2.3 Strawman Solutions

A straightforward solution is only using big cores. However, little cores can help achieve higher throughput under a lower contention (e.g., 68% in Figure 8g). Finding the optimal number of cores to run applications is a long-existing problem [35, 44]. Directly blocking little cores under high contention will cause a latency collapse (Implication 2). Moreover, since the energy-aware scheduler may schedule threads to little cores for energy purposes, binding or migrating threads to big cores will violate the energy target.

Another intuitive solution is adopting proportional execution that gives big cores a fixed higher chance to lock. Figure 5 shows the performance when setting different proportions. The throughput and the latency are mutually exclusive in AMP. A larger proportion brings higher throughput but longer latency. However, it is unclear whether a specific application prefers throughput over latency (and the extent) or the opposite. Moreover, since applications’ loads may change over time, the latency is unstable when setting a fixed proportion (not a problem in SMP since the fairness

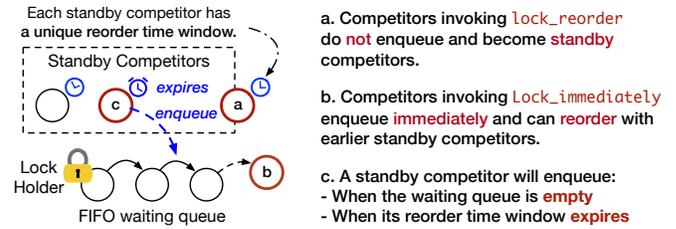


Figure 7. Reorderable lock blocks standby competitors and allows other competitors to reorder with them.

always ensures the lowest latency). Therefore, it is almost impossible to find a suitable proportion to meet various applications’ needs.

## 3 Design of LibASL

### 3.1 Overview

To address the lock scalability problem on AMP, we propose an asymmetry-aware scalable lock LibASL. Rather than preserving the lock acquisition fairness, LibASL provides a new lock ordering guided directly by the applications’ latency requirements (i.e., SLO) for better throughput (according to Implication 1). Atop of a FIFO waiting queue, LibASL allows reordering under the condition that the victim (reordered) will not miss the application’s latency SLO. Thus, big cores can reorder as much as possible with little cores to achieve higher throughput, while little cores can barely meet their latency SLO (according to Implication 2).

To achieve the SLO-guided ordering, bounded reordering is needed. Thus, we first design a *reorderable lock*, which exposes the bounded reorder capability as a configurable reorder time window atop of a FIFO waiting queue. Only during the time window, big cores can reorder (lock before) with little cores. Once the window expires, no reorder will happen (bounded). However, it is non-trivial to set a suitable fine-grained window for each lock acquisition based on the application’s coarse-grained latency requirement. To this end, by proposing a feedback mechanism, LibASL automatically chooses a suitable reorder window on each lock acquisition according to the coarse-grained latency SLO.

**Usage model.** LibASL is easy-to-use. As shown in Figure 6, LibASL provides **two intuitive interfaces** to annotate the latency SLO of a certain code block (named as an *epoch*), including `epoch_start` and `epoch_end`. Each epoch has a unique epoch id, which is statically given by programmers and should be passed as an argument (e.g., 5 in Figure 6). `epoch_end` takes another argument, which specifies the latency SLO of the epoch in nanoseconds (e.g., 1000 means the epoch’s latency SLO is 1µs in Figure 6). LibASL does not restrict either the number of locks or how the lock is used in an epoch. Thus, programmers can directly mark a coarse-grained latency SLO without dividing (e.g., the whole request handler in Figure 6) and create as many epochs as

needed. LibASL leverages weak-symbol replacement to redirect the invocations of `pthread_mutex_lock` transparently. Therefore, no other modification is required.

LibASL can be used in various applications to improve the scalability under AMP with minimal effort.

- For latency-critical applications, annotating their existing coarse-grained SLOs is the **only** required effort to use LibASL. No other knowledge is required. Such SLOs are defined according to actual latency targets and are commonly available in practice. For example, an interactive app has an SLO of 16.6ms to satisfy the 60Hz frame-rate requirement.
- For applications without clear SLOs, LibASL provides a profiling tool that generates a latency-throughput graph (e.g., Figure 8b) to help choose suitable SLOs. The graph is generated by automatically iterating different SLO settings inside a given SLO range.
- Non-latency-critical applications can transparently use LibASL with no SLO. LibASL directly uses a default (fairly loose) reorder time window to maximize the throughput without causing starvation.

### 3.2 Reorderable Lock

The reorderable lock exposes the bounded reorder capability atop of existing FIFO locks (e.g., the MCS lock). It provides two interfaces to acquire the lock, including `lock_reorder` and `lock_immediately`. Figure 7 presents the behavior of both interfaces. Competitors invoking `lock_immediately` will be appended to the tail of the waiting queue immediately. Competitors invoking `lock_reorder` will be regarded as standby competitors. If the waiting queue is empty, standby competitors can enqueue and then become the lock holder. Otherwise, the standby competitors are blocked. Each standby competitor has a unique reorder time window, which is an argument of `lock_reorder`. Other competitors can reorder with them and lock earlier during that window. Thus, the reordering is bounded by that window. A standby competitor will enqueue once its reorder window expires.

Algorithm 1 shows the implementation of the reorderable lock. When calling `lock_immediately`, the competitor will directly enqueue (line 2) by using the lock interface (`lock_fifo`) of the underneath replaceable FIFO lock (e.g., MCS). `lock_reorder` takes an argument window that specifies the length of the reorder window in nanoseconds. When calling `lock_reorder`, the competitor will firstly check whether the lock is free (line 7). If so, the competitor will enqueue immediately. Otherwise, it will become a standby competitor. During the reorder window, the standby competitors will occasionally check the lock's status (line 11). We use a binary exponential back-off strategy (line 12)

**Algorithm 1.** Reorderable lock implementation.

```

1 int lock_immediately(mutex_t *mutex) {
2     return lock_fifo(mutex);
3 }
4
5 int lock_reorder(mutex_t *mutex, u64 window) {
6     u64 window_end, cnt = 0, next_check = 1;
7     if (is_lock_free(mutex)) goto out;
8     window_end = current() + window;
9     while (current() < window_end) {
10        if (cnt ++ == next_check) {
11            if (is_lock_free(mutex)) goto out;
12            next_check <<= 1;
13        }
14    }
15 out:
16    return lock_fifo(mutex);
17 }
18
19 int unlock(mutex_t *mutex) {
20    return unlock_fifo(mutex);
21 }

```

to reduce the contention over the lock<sup>3</sup>. When the reorder time window expires, the competitor can finally enqueue. We do not use a secondary queue for the standby competitors because each competitor can have a different reorder window. Thus, they may enqueue at a different time once the reorder window expires (not in FIFO order). We also set an upper bound for the reorder window (omitted here) to make the reorderable lock starvation-free. The reorder window is **not** a strict order constraint. It is only a hint about the time the competitor can wait considering its latency SLO. Thus, although big cores can still lock first when it expires (i.e., after invoking `lock_fifo` in line 16 but have not enqueued yet), it does not influence its correctness or efficiency.

Since the reorderable lock does not modify the underneath lock, for unlocking, the reorderable lock directly invokes the unmodified unlock procedure `unlock_fifo` (line 20).

### 3.3 LibASL

Atop of the reorderable lock, LibASL collects the application's latency SLO and chooses a suitable reorder window accordingly to maximize the reordering without violating the SLO. The mapping is achieved by tracing all epochs' latency and adjusting the reorder window at every epoch ends. LibASL keeps individual reorder windows for each epoch. When locking in an epoch, LibASL calls `lock_immediately` if running on big cores. Otherwise, it calls `lock_reorder` and sets the reorder window of that epoch.

Algorithm 2 shows the implementation of the LibASL's epoch interfaces. Each epoch has the per-thread metadata, which keeps the length of the reorder window (`window`), the start timestamp (`start`) and the adjusting unit of the epoch's length (`unit`). When initializing, we give a default size to both the window and `unit`. They will quickly adjust themselves to a suitable size after executing a few epochs.

<sup>3</sup>We present the non-blocking (busy-waiting) implementation here. LibASL also has a blocking version that yields the thread using `nanosleep` during the reorder time window. We evaluate both versions in Section 4.

**Algorithm 2.** LibASL epoch implementation.

```

1 typedef struct epoch {
2   u64 window; /* Reorder Window*/
3   u64 start; /* Timestamp */
4   u64 unit; /* Adjust Unit */
5 } epoch_t;
6 __thread epoch_t epoch[MAX_EPOCH];
7 __thread int cur_epoch_id = -1;
8 __thread int *epoch_stack;
9 #define PCT 99 /* 99th Percentile Latency */
10
11 int epoch_start(int epoch_id) {
12   if (cur_epoch_id >= 0)
13     push(epoch_stack, cur_epoch_id);
14   cur_epoch_id = epoch_id;
15   epoch[epoch_id].start = current();
16   return 0;
17 }
18
19 int epoch_end(int epoch_id, u64 SLO) {
20   u64 latency, window;
21   if (is_big_core()) goto out;
22   latency = current() - epoch[epoch_id].start;
23   window = epoch[epoch_id].window;
24   if (latency > SLO) {
25     window >>= 1;
26     epoch[epoch_id].unit = window * (100 - PCT) / 100;
27   } else {
28     window += epoch[epoch_id].unit;
29   }
30   epoch[epoch_id].window = window;
31 out:
32   cur_epoch_id = empty(epoch_stack) ? \
33     -1 : pop(epoch_stack);
34   return 0;
35 }

```

When calling `epoch_start`, `epoch_id` specifies the unique id of the upcoming epoch, which will be stored in the per-thread global variable `cur_epoch_id` (line 14). To support the nested epoch, it pushes the outer epoch to the stack if it exists (line 13). Then it records the start timestamp (line 15) by using the light-weight `clock_gettime` (~45 cycles).

Besides the `epoch_id`, `epoch_end` takes another argument SLO, which specifies the latency SLO of the current epoch in nanoseconds. It calculates the latency (line 22), compares it with the SLO (line 24) and updates the reorder window accordingly. We take a conservative strategy to adjust the reorder window inspired by the TCP congestion control algorithm [25], which combines linear growth and exponential reduction when latency exceeds. We set the granularity of growth (unit) to be  $\frac{100-PCT}{100}$  of the reduced window<sup>4</sup>, where PCT represents the percentile the SLO specifies (line 9, other percentiles are also supported). It then checks the epoch stack to see whether there is a nested epoch and sets `cur_epoch_id` accordingly (line 32-33).

By leveraging weak-symbol replacement, LibASL redirects `pthread_mutex_lock` in applications to `asl_mutex_lock` in Algorithm 3 transparently with negligible overhead (20+ cycles, similar to `litl` [44]). When calling `libASL_lock`, competitors from big cores directly acquire the underneath lock using `lock_immediately` (line 3), while those from little

<sup>4</sup>After another  $\frac{100}{100-PCT}$  executions, the latency will be the same as the one which barely exceeds the SLO and triggers the exponential reduction. The probability of not exceeding the SLO is  $(\frac{100}{100-PCT} - 1) / \frac{100}{100-PCT} = \frac{PCT}{100}$ .

**Algorithm 3.** LibASL internal interface.

```

1 int asl_mutex_lock(mutex_t *mutex) {
2   if (is_big_core()) /* Big core */
3     return lock_immediately(mutex);
4   else if (cur_epoch_id < 0) /* Not in epoch */
5     return lock_reorder(mutex, MAX_WINDOW);
6   else /* In an epoch */
7     return lock_reorder(mutex,
8       epoch[cur_epoch_id].window);
9 }

```

cores use `lock_reorder` and set the window length according to the current epoch (line 7-8). If not in any epoch, the default maximum window is used to ensure that the thread will eventually lock (line 5). Identifying the core type is achieved by getting the core id and looking up a pre-defined table. Since the reorderable lock is implemented atop of existing locks, both the trylock and the nested locking are supported. Besides, the conditional variable is also supported by using the same technique in `litl` [44].

### 3.4 Analysis

**Throughput.** LibASL provides good scalability in AMP. We analyze different situations applications may encounter and the corresponding behavior of LibASL as follows.

*Big cores and little cores are **not** competing for the same lock.* In big cores, LibASL behaves the same as the underneath FIFO lock (e.g., the MCS lock). In little cores, LibASL behaves similarly to the backoff spinlock. Both locks are scalable [26] when competitors are from the same type of core.

*Big cores and little cores are competing for the same lock.* When the lock is not heavily contented, competitors from both big and little cores can immediately hold the lock if the lock is free (no additional overhead). Little cores can help achieve higher throughput in such cases (Figure 8g shows the corresponding experiment). With the contention level increased, big cores will reorder with little cores under the condition that the latency SLO is still met. When the big cores do not saturate the lock (i.e., the lock becomes free sometimes in a while), little cores will lock once the queue is empty. Thus, LibASL can find the sweet spot where some additional little cores help saturate the lock for better throughput (and block the rest little cores). Otherwise, allowing any extra little core to join the competition will degrade the throughput. In those cases, little cores can get the lock only when the reorder window expires. Thus, LibASL can improve the throughput as much as the latency SLO allows.

**Latency.** LibASL can precisely maintain the latency under SLO through a feedback mechanism. The size of the reorder window has a monotonic relationship with the epoch's latency (i.e., a smaller window means a shorter waiting time). It still holds when an epoch contains multiple lock acquisitions since they share the same window size. Thus, LibASL can find the suitable window size that the latency barely meets the SLO by adjusting the size according to the latency (if the latency is higher than SLO, shrink the window, and

vice versa). Even if the epoch length (i.e., execution time) becomes heterogeneous (e.g., executing different code paths), LibASL can still maintain the tail latency because the reorder window shrinks exponentially once the violation happens and grows linearly in the subsequent executions. Thus, it only gives some short epoch a small reorder window (larger window can still meet SLO) but will not violate the SLO.

LibASL also supports threads migration or co-running with other applications. Since each epoch’s metadata is per thread, they will not influence each other inherently. Normally, the lock’s contention remains stable in a while. Thus, no extra window adjustment is required. Otherwise, when the thread gets scheduled out inside an epoch, or the lock’s contentions vary significantly, the window will quickly adjust itself in the same way as facing heterogeneous epochs.

Notice that the latency SLO is not a strict deadline. LibASL uses it as a hint to maximize throughput without violating it. There are three cases where LibASL does not take effect. First, when the SLO is impossible to achieve even without reordering, LibASL falls back to a FIFO lock (best effort). Second, when the SLO of nested epochs are mistakenly set (e.g., outer epoch has a tighter SLO), LibASL always prioritizes the inner epoch. Third, if the workload is non-lock-sensitive (locks are barely contended), LibASL inherently will not influence its performance.

**Energy.** Energy-efficiency is one of the major targets of the AMP system. To maintain the lowest energy consumption, Linux provides EAS (energy-aware scheduler [9]), which chooses the most suitable core for each thread. LibASL does not require core-binding, and threads can migrate between cores freely. Therefore, LibASL will not violate the scheduling decision, as well as the energy target. Threads will only run on big cores when the scheduler decides so (not caused by LibASL). Moreover, when running on big cores, LibASL makes threads do meaningful jobs rather than waiting, which helps to save energy [39].

**Target systems.** LibASL can be used in various AMPs. Its improvement comes from considering the asymmetric computing capacity, and it does not restrict to a certain AMP. Future AMPs may also have large core counts or even NUMA. LibASL can adapt to those AMPs by replacing the underlying lock with the corresponding scalable locks (e.g., NUMA-aware locks). It will prioritize big cores in the upper reorderable lock while achieving good scalability in the waiting queue of the underlying lock (e.g., NUMA-locality).

Some AMPs also support DVFS. Although LibASL does not explicitly consider DVFS, in most AMPs, a big core, even with the lowest OPP (Operating Performance Points), still has better performance than the highest little core OPP [70]. Thus, LibASL still brings improvement by prioritizing faster big cores. For other platforms (i.e., big cores may sometimes be slower than little cores), LibASL requires an extra mechanism to be aware of the performance variation due to DVFS.

**Limitations.** LibASL brings improvement when the application is lock-sensitive and has a relatively loose SLO. Otherwise, when the SLO is too tight (only achievable when passing the lock in FIFO), or the workload is non-lock-sensitive, LibASL behaves the same as the underneath FIFO lock.

LibASL itself does not bring noticeable space or performance overhead. The space overhead of LibASL comes from the metadata of epochs, which is negligible because the per-thread metadata of an epoch only takes 24 bytes (see Algorithm 2) and is irrelevant to the number of locks; the two epoch operations only involve cheap computations ( $\sim 93$  cycles per epoch), which barely influence the performance.

## 4 Evaluation

We evaluate LibASL to answer the following questions:

1. How much throughput can LibASL improve when setting variant SLOs?
2. Can LibASL precisely maintain epochs’ latency in various situations?
3. How does LibASL perform under variant contention levels?
4. Can LibASL take effect in real-world applications?

**Evaluation Setup.** We evaluate LibASL on Apple M1, the only available desktop AMP yet. M1 has 4 big cores and 4 little cores. Similar to other AMP [52, 66], their performances gap varies in different applications. Big cores are 3.75x faster in Sysbench [21], while only 1.8x faster when executing the NOP instruction. In such a system, the theoretical throughput speedup upper bound of LibASL to a FIFO lock (e.g., MCS) by prioritizing big cores is  $1.8x^5$ . We run Linux 5.11 on M1 [8].

Unless explicit statements, the reorderable lock is built atop the MCS lock, and the PCT is set to 99 (P99 latency). We bind threads to different cores to evenly distribute them for stable results (not required by LibASL), which is a widely-adopted evaluation method [35, 36, 38, 42, 47, 50, 51, 58, 59, 63, 64, 71]. We compare LibASL with `pthread_mutex_lock` (glibc-2.32), TAS, ticket, MCS and `ShflLock` [50]. `ShflLock` provides a lock reordering framework but can only take a static policy (e.g., NUMA-local policy prioritizes competitors from the same node). The SLO-guided ordering in LibASL is not static. Thus it is hard to be integrated. Instead, we adopt a proportional-based static policy, which gives the big core a fixed higher (10x) chance to lock. It is implemented by modifying the existing NUMA-local policy to separate asymmetric cores into two nodes. Rather than passing in each node evenly (for preserving long-term fairness in NUMA-local policies), it uses a simple counter to allow exactly 1 little core to lock after every 10 big cores. Since any proportion is one static trade-off between latency and throughput

<sup>5</sup>  $\frac{4.75+1}{2} - 1 \approx 1.8$ : Comparing the case where the big cores always run and the case where the big cores and little cores run one by one.

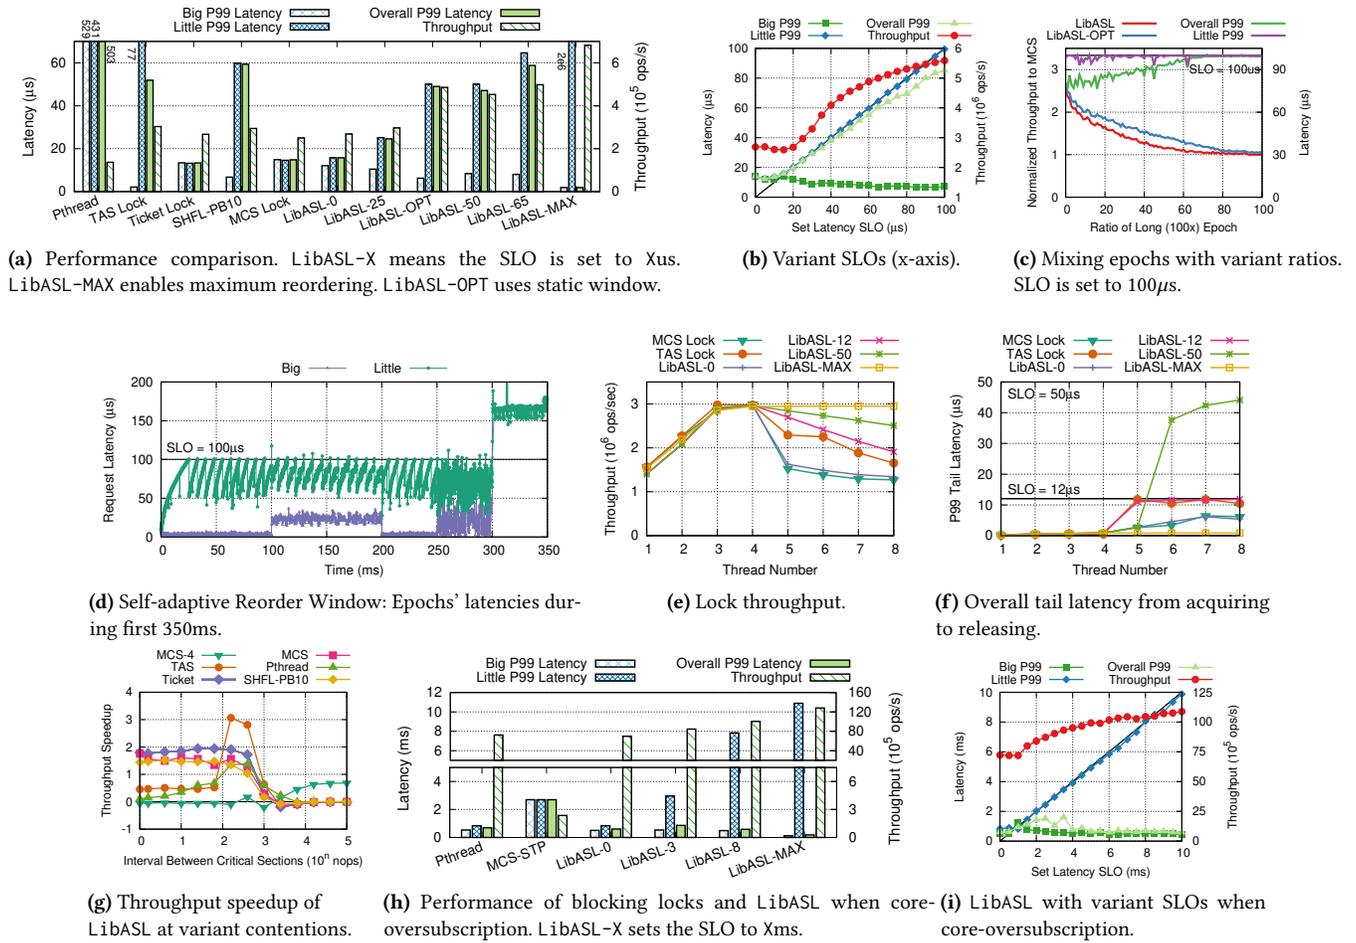


Figure 8. Micro-benchmarks. Big P99 and Little P99 presents the 99th percentile latency in big and little cores individually.

(Figure 5), we choose the proportion (i.e., 10) that has obvious throughput improvement without introducing extremely long latency. We also present the speedup upper bound of LibASL by using the maximum reorder window (100ms).

### 4.1 Micro-Benchmarks

**Bench-1: A heavily contended benchmark.** In this benchmark, all threads repeatedly execute the same epoch, which contains 4 critical sections of different lengths protected by 2 different locks. In each critical section, threads read-modify-write a specific number of shared cache lines (64 in total). We insert a fixed number ( $600 \times 2^7$ ) of NOP instructions between two epochs. We present epoch's tail latency in little cores, big cores and overall separately.

As shown in Figure 8a, the TAS lock shows big-core-affinity here (big cores have a shorter tail latency). Thus, it has the highest throughput yet the longest latency among existing locks. When setting the SLO to 0 (LibASL-0), LibASL performs the same as the MCS lock since the SLO is impossible to achieve (falls back to FIFO). Compared with the TAS lock, when achieving a similar throughput (LibASL-25),

LibASL reduces the tail latency by over 50%; when having a similar tail latency (LibASL-50), LibASL achieves 50% better throughput. Although TAS lock also implicitly prioritizes the big cores here, the reordering depends on the hardware and is unstable and uncontrollable. LibASL manages the reordering elaborately, which allows more critical sections to execute on the big cores and thus has higher throughput. LibASL brings up to 1.2x speedup to TAS lock (1.7x to MCS) when setting a larger SLO (LibASL-MAX). Although the proportional-based approach (SHFL-PB10) performs better than MCS by 20%, it has a 4x longer tail latency. LibASL outperforms it by 70% when having similar tail latency (LibASL-65). It is because LibASL has a better cache locality by batching more big cores before passing to little cores for no SLO violation. In contrast, the proportional-based approach has to periodically (i.e., every 10 big cores) hand over the lock to little cores. The pthread\_mutex\_lock has the worst throughput and the longest latency. LibASL outperforms it by 4x at most. **(Question 1)**

We also compare LibASL with the optimal policy LibASL-OPT, which directly chooses a static window (no

window adjustment) to present LibASL’s performance headroom. When having a similar tail latency (LibASL-50), the cost of the window adjustment is only 6%.

Figure 8b shows LibASL’s performance in *Bench-1* when setting variant SLOs. As shown in the figure, with setting a larger SLO (x-axis from left to right), the throughput increases, and the tail latency of little cores sticks straightly to the  $Y=X$  line (i.e., barely meet the SLO). Meanwhile, since big cores get more chances to lock, their latencies are much shorter. The growth speed of throughput slows down with the SLO becoming larger. It is intuitive because the benefit of reordering more big cores will decrease if most critical sections are already executed on the faster big cores. The only exception is when setting an SLO shorter than  $15\mu\text{s}$  (the tail latency of MCS in Figure 8a), the SLO is impossible to achieve. Therefore, LibASL falls back to the MCS lock.

**Bench-2:** *A highly variable workload.* We record each epoch’s latency in the first 350ms when executing *Bench-1*. Figure 8d shows the latency of each epoch executed on big and little cores individually. The SLO is set to  $100\mu\text{s}$ . During the 100 and 200ms period, we enlarge the epoch’s length by 128 times (by accessing more cache lines) and shrink it back to the original length during the 200 and 250ms period. After that, we change the length of each epoch randomly (access a random number of cache lines) during the 250 to 300ms period. Finally, the epoch’s length is set 1024 times longer till the end. As shown in the figure, LibASL is fully capable of maintaining the latency in a highly variable workload (**Question 2**). Every time the latency exceeds the SLO, the reorder window shrinks to its half and increases gradually. When the epoch’s length changes at 100 ms and 200 ms, LibASL quickly adjusts the reorder window to a suitable size. Even when the length becomes highly heterogeneous during 250 and 300 ms, LibASL can still keep the latency within the SLO. When the epoch’s length enlarges 1024 times at 300 ms, the SLO is impossible to achieve. Thus, LibASL falls back to FIFO, and both big and little cores have similar latencies.

**Bench-3:** *A benchmark mixing epoch of significantly different lengths.* In Figure 8c, we randomly generate short and long ( $100\times$  longer by inserting more NOP instructions) epochs of different ratios (e.g.,  $x=20$  means 20% of epochs are short while 80% are long). We compare LibASL with LibASL-OPT, which directly chooses a suitable (static) window for different epochs (impossible in the real world). The SLO is set to  $100\mu\text{s}$  throughout the experiment. As shown in the figure, LibASL brings significant and close-to-optimal (maximum 20% gap with LibASL-OPT at 50% ratio) throughput improvement to MCS while precisely keeping the latency within SLO in all ratios (**Question 2**). When all epochs are long (i.e.,  $x=100$ ), the tail latency of the MCS lock is also  $100\mu\text{s}$  (the same as the SLO). Thus, LibASL falls back to FIFO for not violating SLO and has the same throughput (i.e.,  $y=1$ ) as the MCS lock.

**Table 1.** Databases Considered

Application	Benchmark	Locks in each Epoch
<b>Kyoto Cabinet</b> [14]	50% Put 50% Get	Slot-level Lock
In-memory KV		Method Lock
<b>upscaledb</b> [22]	50% Put 50% Get	Global Lock
On-disk KV		Worker Pool Lock
<b>LMDB</b> [16, 31]	50% Put 50% Get	Global Lock
On-disk KV		Metadata Locks
<b>LevelDB</b> [15]	db_bench Random Read	Metadata Lock
On-disk KV		
<b>SQLite</b> [20]	1/3 Insert 1/3 Simple Select	State Machine Lock
On-disk Database	1/3 Complex Select	Metadata Locks

**Bench-4:** *Scalability.* We present the scalability of LibASL using the same benchmark setup as Figure 4. As shown in Figure 8e and 8f, when setting the SLO to 0, LibASL-0 behaves the same as the MCS lock. When setting the SLO to have the same tail latency ( $12\mu\text{s}$ ) as the TAS lock, LibASL-12 achieves better throughput scalability. The throughput of LibASL-MAX does not drop at all. Due to the high contention, it barely passes the lock to the little cores.

**Bench-5:** *A benchmark with variant contentions.* In this benchmark, threads acquire the same lock to read-modify-write 2 shared cache lines. We alter the contention by executing a different number of NOP instructions between two lock acquisitions. Figure 8g shows the throughput speedup of LibASL over the locks in the legends (e.g., when  $x=0$ , LibASL outperforms MCS by 2x and TAS lock by 45%). To allow the maximum reordering, we do not set SLO in LibASL. We also include the result of only using big cores (only MCS-4, other locks use all the cores). When competitors from big cores already saturate the lock ( $x < 3$ ), LibASL makes the competitors standby and achieves similar throughput with MCS-4 (significantly better than others). With the contention level decreased, LibASL allows little cores to join the competition and achieves 68% better throughput than only using big cores. It also validates that little cores can bring noticeable improvement in some loads. Among all contention levels, LibASL achieves good throughput. (**Question 3**)

**Bench-6:** *A benchmark with CPU core over-subscription.* We examine the effectiveness of LibASL in a core over-subscription situation by creating 2 threads on each core and executing *Bench-1*. We replace the non-blocking MCS lock in LibASL with the `pthread_mutex_lock` and use `nanosleep` to replace the busy waiting of the standby competitor (i.e., line 9 in Algorithm 1, the sleep time is set in a back-off manner). Results are presented in Figure 8h and 8i. Since the MCS lock passes the lock in a FIFO order, the waking-up latency will be introduced to the critical path, leading to a significant throughput degradation (spin-then-park MCS is 96% worse than `pthread_mutex_lock`). Therefore, LibASL uses `pthread_mutex_lock` rather than the spin-then-park MCS lock. Although `pthread_mutex_lock` does not guarantee the FIFO order and thus has an unstable lock acquisition latency,

LibASL still can preserve the SLO owing to its self-adaptive reorder window and outperform the pthread\_mutex\_lock by up to 80%.

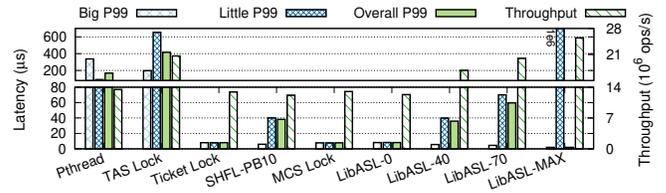
### 4.2 Application Benchmarks

We evaluate 5 popular databases in Table 1 to show the effectiveness of LibASL in the real-world (Question 4). Databases benefit from using little cores to handle more requests in fewer machines, which can improve the cost and energy efficiency in edge computing [57, 68, 72]. Integrating LibASL only requires inserting 3 lines of code: annotating the epoch with epoch\_start and epoch\_end and adding the header file. As prior work [36, 50] does, we run each benchmark for a fixed period and calculate the average throughput. Moreover, to present the effectiveness of LibASL in cases where epochs' lengths are highly heterogenous, we randomly choose to insert or find 1 item (fifty-fifty, referring to YCSB-A [23]) in an epoch. In most benchmarks, each epoch acquires multiple locks as listed in the rightmost column of Table 1. For each database, we first set several specific SLOs to present a performance comparison with existing locks to show the improvement of LibASL when having similar latency or throughput with existing locks (e.g., TAS). Then we show the performance under other SLOs in Variant SLOs figures.

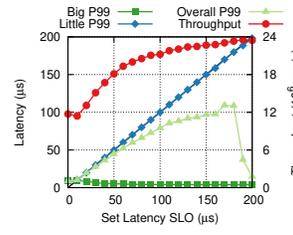
We first evaluate LibASL in several KV-stores. KV-store plays an important role in CDN or IoT edge servers as the storage service [1, 6]. As shown in Figure 9a, in Kyoto Cabinet, LibASL reduces the tail latency by 90% when achieving similar throughput with the TAS lock (LibASL-70) and improves the throughput by up to 23% with a larger SLO (96% to MCS, 89% to pthread\_mutex\_lock). When having a similar latency with SHFL-PB10 (LibASL-40), LibASL improves the throughput by 38%. Figure 9b shows the performance of LibASL when setting variant SLOs. Although the execution time of Put or Get is heterogeneous, LibASL can still precisely maintain the tail latency.

Figure 9c presents the latency Cumulative Distribution Function (CDF) of LibASL when setting the SLO to 70μs. Overall and Little represent the overall and little core's latency. A clear boundary can be seen in the overall latency since most operations are executed on big cores. Due to the intensive contention, only less than 20% of operations are executed on little cores and have longer latency. There is also a clear boundary in little core's latency. About half of the operations have shorter latency (<35μs) due to the shorter execution time of Get. As for the longer Put operation (another half), since the reorder window shrinks by half and grows linearly once the latency exceeds the SLO, the probability also grows linearly after the Half SLO (35μs).

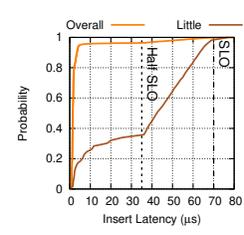
Upscaledb and LMDB show similar results (Figure 9d~9i). In upscaledb, the TAS lock shows big-core affinity. Thus it has 90% higher throughput yet 2.5x longer tail latency than the MCS lock. When having a similar tail latency (LibASL-140), LibASL improves the throughput of the TAS lock by 46%,



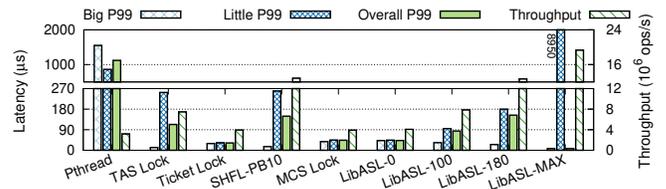
(a) Kyoto Cabinet. 1e6 means  $10^6$ . The chosen SLOs are only for easier comparing. Other settings are detailed in figure 9b.



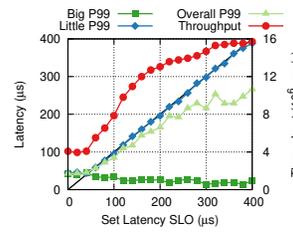
(b) Variant SLOs



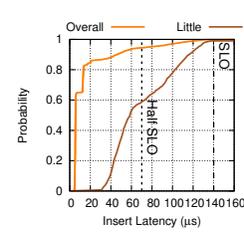
(c) CDF (SLO: 70μs)



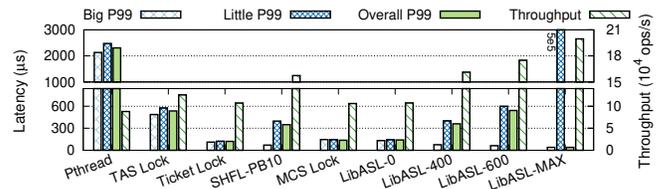
(d) upscaledb



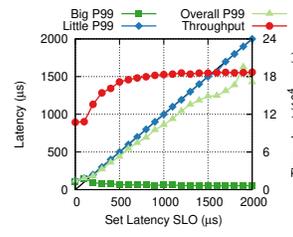
(e) Variant SLOs



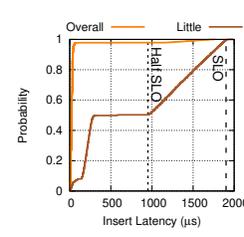
(f) CDF (SLO: 140μs)



(g) LMDB. 5e5 means  $5 \times 10^5$ .



(h) Variant SLOs



(i) CDF (SLO: 1900μs)

Figure 9. Databases. Legends are explained in Figure 8.

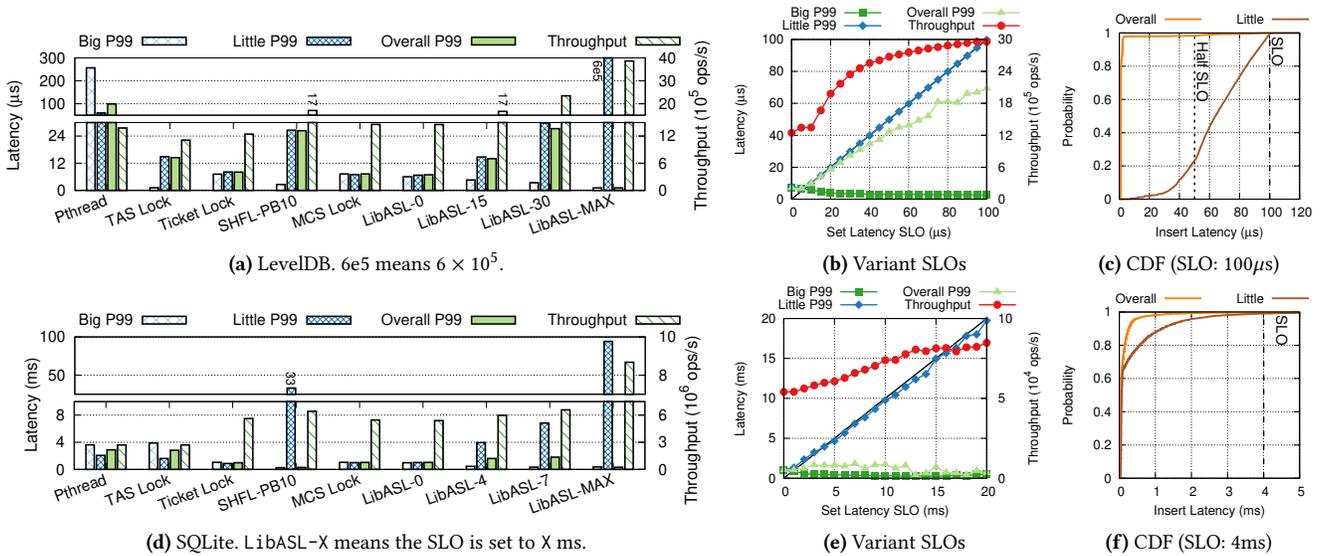


Figure 10. Database benchmarks. Legends are explained in Figure 8.

which further goes to 1.6x (3.8x to MCS, 50% to SHFL-PB10 and 5x to pthread\_mutex\_lock) with a larger SLO. In LMDB, LibASL outperforms the TAS lock by 40% when having a similar latency (LibASL-600), which goes to 60% (86% to MCS, 27% to SHFL-PB10 and 126% to pthread\_mutex\_lock) with a larger SLO. In the latency CDFs of both benchmarks (Figure 9f, 9i), the clear boundary in little core’s latency distinguishes the shorter Get from the longer Put operation.

LevelDB is another widely-used KV-store. However, LevelDB implements its own blocking strategy rather than directly using the pthread\_mutex\_lock on Put operation. Thus we use the randomread test in the build-in db\_bench to only test its Get operation, which will acquire a global lock to take a snapshot of internal database structures. As shown in Figure 10a, LibASL improves the throughput of TAS lock by 50% when having a similar latency (LibASL-15), which goes to 2.5x (1.6x to MCS, 1.8x to pthread\_mutex\_lock and 1.3x to SHFL-PB10) with a larger SLO. Since we only test the Get operation, most requests have a longer latency than the *Half SLO* as shown in Figure 10c.

Finally, we evaluate LibASL in SQLite, which has been used in Azure IoT edge servers [5]. We place 1/3 Insert, 1/3 simple query (point query on an indexed column) and 1/3 complex query (range query on an indexed column with a filter on a non-indexed column) in a DEFERRED transaction enclosed in an epoch. Moreover, we add an extremely long full-table scan on a 100k table every 1000 executions in the same epoch to present that LibASL can survive on some occasionally appeared extremely long requests. SQLite uses locks to protect the internal state machine. The transaction can commit successfully only in a certain state. Thus epochs’ latencies greatly fluctuate and grow non-linearly in Figure 10f. It also widens the gap of the transaction’s success rate in big and little cores when using SHFL-PB10 and causes

a latency collapse in little cores. Both the simple and the complex Select operations have a much shorter execution time than the Insert operation. Thus, 2/3 of the requests have shorter tail latency (latency grows significantly after  $y > 2/3$  in Figure 10f). However, even with some occasionally appeared extremely long epochs, LibASL still can precisely keep the tail latency under the SLO and improve the throughput as shown in Figure 10e. LibASL brings up to 2.1x speedup to the TAS lock (55% to MCS, 2.1x to pthread\_mutex\_lock and 35% to SHFL-PB10) without violating SLO.

Besides M1, we also evaluated LibASL in Hikey970 [11] (ARM big.LITTLE) and a simulated Intel AMP (through per-core DVFS). LibASL works well on both platforms since the improvement comes from considering the asymmetry in computing capacity and is **not** restricted to a certain AMP. Specifically, LibASL brings 34~94% (Intel) and 37~87% (Hikey970) throughput improvement to the MCS lock while precisely maintaining the SLO in the same database benchmarks. Detailed results are omitted due to the space limit.

### 4.3 Evaluation Highlights

Results confirm the effectiveness of LibASL in AMP. First, LibASL can precisely maintain the latency SLO even in highly variant workloads. Second, LibASL shows promising performance advantages over existing locks. Compared to fair locks (lowest latency but low throughput), LibASL can significantly improve the throughput (e.g., 3.8x to the MCS lock in upscaledb). Compared to unfair locks (highest latency but sometimes high throughput), LibASL has much lower tail latency when achieving similar throughput (e.g., 90% lower than the TAS lock in KyotoCabinet), and substantially higher throughput when ensuring similar tail latency (e.g., 46% higher than the TAS lock in LMDB). Moreover, LibASL can further outperform the TAS lock when setting

a larger SLO (e.g., by 2.5x in LevelDB). Compared to the static proportional-based approach, LibASL can better meet applications' needs by improving the throughput as much as possible considering a certain latency SLO that a fixed proportion cannot.

## 5 Related Work

Scalable synchronization primitives [28–30, 32, 35–38, 41–43, 45, 47, 49, 50, 55, 56, 58–60, 62–65, 71] have been extensively studied over decades, targeting various scenarios. Yet, there lacks investigation on the scalability problem on AMP.

There are already some locks [29, 32, 35–38, 50, 59, 63] that reorder competitors to achieve better throughput in NUMA or many-core processors, which inspire LibASL. Among those locks, defining an effective lock ordering (i.e., how to reorder) to achieve different optimization targets is important and challenging. The major difference between LibASL and those locks is that LibASL defines a new SLO-guided AMP-specific ordering while existing lock orderings (e.g., NUMA-local) are non-scalable in AMP. As analyzed in Section 2.2, those locks preserve the long-term fairness to keep a relatively low latency, which brings throughput collapse on AMP. Besides, unlike LibASL, the long-term fairness only forbids starvation, leaving the latency unpredictable. Rather than proposing a specific ordering, ShflLock [50] provides a reordering framework. It relies on a provided *static* policy to shuffle the queue internally. However, the scalability issue in AMP cannot be easily solved by proposing a static policy. Preserving fairness brings throughput collapse; reordering without limit causes latency collapse; proportional execution brings unstable and unpredictable performance so that a suitable static proportion is hard to find as evaluated in Section 4. Instead, LibASL has a *dynamic* priority to achieve good AMP scalability.

Delegation locks [41, 42, 47, 58, 62, 64, 71] reduce the data movement by executing all critical sections in one core (the lock server), which significantly improves the throughput in NUMA. Although placing the lock server on big cores can hide the weak computing capacity of little cores, it requires the big core busy polling, which wastes a precious big core and violates the energy target at lower contention. A more severe obstacle to using the delegation lock is that it requires non-trivial code modifications to convert all critical sections into closures, which brings enormous engineering work due to the complexity of real-world applications. Instead, LibASL only requires linking and, if latency-critical, inserting few lines of code to specify the latency SLO.

There are also self-tuning locks [45, 49, 60] that tune some internal parameters at runtime. LibASL can be regarded as one such lock considering its self-adaptive reorder window. However, existing self-tuning locks target solving different problems and thus have different designs. Specifically, reactive lock [45] tunes the back-off time in spinlock to reduce

the contention; mutlock [60] tunes the number of busy spinning threads to hide the wakeup latency in blocking locks (also named as *window*, but it is a different mechanism: its *spinning window* is per-lock and restricts how many threads can busy spin; LibASL's *reorder window* is per-thread and controls how long others can reorder); locks in [49] tune the spinning threshold before parking for better throughput. Unlike them, LibASL tunes the reorder window to preserve a new AMP-aware lock ordering for better scalability on AMP. Existing self-tuning locks do not preserve a specific lock ordering and fail to scale on AMP.

LibASL blocks some competitors from little cores to prioritize big cores. Some existing locks [35, 49, 60] also leverage concurrency restrictions for different targets. Mutlock [49] and locks in [49] park some competitors to save CPU; Malthusian lock [35] only allows a subset of competitors to lock for reducing contention. Those locks target SMP, and their techniques cannot solve the scalability issue in AMP. Facing the asymmetry, it is challenging to restrict certain competitors for better performance. LibASL adopts an SLO-guided runtime restriction (through reorder window) and achieves good scalability in AMP.

Computing capacity can also be asymmetric in SMP when using DVFS. Previous work [24, 27, 67] boosts the frequency of the lock holder to gain better throughput. However, unlike DVFS, the asymmetry in AMP is inherent. Thus, those techniques also cannot be applied to AMP.

Improving the system's throughput for better cost and energy efficiency without violating the latency SLO is a widely adopted technique [68, 69, 72]. WorkloadCompactor [72] and Cake [68] reduce the datacenter's cost by consolidating more loads into fewer servers without violating the latency SLO. Elfen Scheduling [69] leverages the SMT to run latency-critical and other requests simultaneously to improve the utilization without compromising the SLO. LibASL takes a similar approach to solve the lock scalability issue in AMP.

## 6 Conclusion

In this paper, we propose an asymmetry-aware scalable lock named LibASL. It provides a new latency SLO-guided lock ordering to prioritize big cores for better throughput while elaborately maintaining little cores' latencies. Evaluations on real-world applications show that LibASL brings better performance over counterparts on AMP.

## 7 Acknowledgement

We sincerely thank all the anonymous reviewers for their insightful suggestions. This work is supported in part by China National Natural Science Foundation (No. 61925206), High-Tech Support Program from Shanghai Committee of Science and Technology (No. 19511121100), and Huawei. Jinyu Gu is the corresponding author.

## References

- [1] [n. d.]. Akamai: IoT Edge Connect. <https://www.akamai.com/cn/zh/products/performance/iot-edge-connect.jsp>.
- [2] [n. d.]. AMD Strix Point Hybrid (Big-Little) CPU. <https://www.hardwaretimes.com/amd-strix-point-hybrid-big-little-cpu-to-feature-3nm-zen-5-cores-zen-4d-cores-l4-cache/>.
- [3] [n. d.]. Apple M1 Chip. <https://www.apple.com/mac/m1/>.
- [4] [n. d.]. ARM DynamIQ Shared Unit Technical Reference Manual. <https://developer.arm.com/documentation/100453/0002/functional-description/introduction/about-the-dsu>.
- [5] [n. d.]. Azure IoT Edge SQLite Module. <https://github.com/Azure/iot-edge-sqlite>.
- [6] [n. d.]. The best IoT Databases for the Edge – an overview and compact guide. <https://objectbox.io/the-best-iot-databases-for-the-edge-an-overview-and-compact-guide/>.
- [7] [n. d.]. CFS wakeup path and Arm big.LITTLE/DynamIQ. <https://lwn.net/Articles/793379/>.
- [8] [n. d.]. CORELLIUM: How We Port Linux to M1. <https://corellium.com/blog/linux-m1>.
- [9] [n. d.]. Energy Aware Scheduling. <https://www.kernel.org/doc/html/latest/scheduler/sched-energy.html>.
- [10] [n. d.]. Google Cloud: Defining SLOs. <https://cloud.google.com/solutions/defining-SLOs>.
- [11] [n. d.]. HiKey970. <https://www.96boards.org/product/hikey970/>.
- [12] [n. d.]. Intel Alder Lake: Performance Hybrid with Golden Cove and Gracemont for 2021, Intel Architecture Day 2020. <https://newsroom.intel.com/press-kits/architecture-day-2020/>.
- [13] [n. d.]. Intel Core i5-L16G7 Processor. <https://ark.intel.com/content/www/us/en/ark/products/202777/intel-core-i5-l16g7-processor-4m-cache-up-to-3-0ghz.html>.
- [14] [n. d.]. Kyoto Cabinet: a straightforward implementation of DBM. <https://dbmx.net/kyotocabinet/>.
- [15] [n. d.]. LevelDB. <https://github.com/google/leveldb>.
- [16] [n. d.]. LMDB TECHNICAL INFORMATION. <https://symas.com/lmdb/technical/>.
- [17] [n. d.]. A Look at Intel Lakefield: A 3D-Stacked Single-ISA Heterogeneous Penta-Core SoC. <https://fuse.wikichip.org/news/3417/a-look-at-intel-lakefield-a-3d-stacked-single-isa-heterogeneous-penta-core-soc/>.
- [18] [n. d.]. New Intel Core Processors with Intel Hybrid Technology. <https://www.intel.com/content/www/us/en/products/docs/processors/core/core-processors-with-hybrid-technology-brief.html>.
- [19] [n. d.]. Processing Architecture for Power Efficiency and Performance. <https://www.arm.com/why-arm/technologies/big-little>.
- [20] [n. d.]. SQLite. <https://www.sqlite.org/index.html>.
- [21] [n. d.]. Sysbench: Scriptable database and system performance benchmark. <https://github.com/akopytov/sysbench>.
- [22] [n. d.]. upscaledb: embedded database technology. <https://upscaledb.com/>.
- [23] [n. d.]. YCSB Core Workloads. <https://github.com/brianfrankcooper/YCSB/wiki/Core-Workloads>.
- [24] S. Akram, J. B. Sartor, and L. Eeckhout. 2016. DVFS performance prediction for managed multithreaded applications. In *2016 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. 12–23. <https://doi.org/10.1109/ISPASS.2016.7482070>
- [25] Mark Allman, Vern Paxson, Wright Stevens, et al. 1999. TCP congestion control. (1999).
- [26] Silas Boyd-Wickizer, M Frans Kaashoek, Robert Morris, and Nikolai Zeldovich. 2012. Non-scalable locks are dangerous. In *Proceedings of the Linux Symposium*. 119–130.
- [27] Juan M. Cebrian, Daniel Sánchez, Juan L. Aragón, and Stefanos Kaxiras. 2013. Efficient inter-core power and thermal balancing for multicore processors. *Computing* 95, 7 (2013), 537–566. <https://doi.org/10.1007/s00607-012-0236-6>
- [28] Milind Chabbi, Michael Fagan, and John Mellor-Crummey. 2015. High Performance Locks for Multi-Level NUMA Systems. *SIGPLAN Not.* 50, 8 (Jan. 2015), 215–226. <https://doi.org/10.1145/2858788.2688503>
- [29] Milind Chabbi and John Mellor-Crummey. 2016. Contention-Conscious, Locality-Preserving Locks. *SIGPLAN Not.* 51, 8, Article 22 (Feb. 2016), 14 pages. <https://doi.org/10.1145/3016078.2851166>
- [30] Haibo Chen, Heng Zhang, Ran Liu, Binyu Zang, and Haibing Guan. 2016. Fast Consensus Using Bounded Staleness for Scalable Read-Mostly Synchronization. *IEEE Trans. Parallel Distributed Syst.* 27, 12 (2016), 3485–3500. <https://doi.org/10.1109/TPDS.2016.2539953>
- [31] Howard Chu. 2011. MDB: A memory-mapped database and backend for OpenLDAP. In *Proceedings of the 3rd International Conference on LDAP, Heidelberg, Germany*. 35.
- [32] Rafael Lourenco de Lima Chehab, Antonio Paolillo, Diogo Behrens, Ming Fu, Hermann Härtig, and Haibo Chen. 2021. CLoF: A Compositional Lock Framework for Multi-Level NUMA Systems. In *Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles (Virtual Event, Germany) (SOSP '21)*. Association for Computing Machinery, New York, NY, USA, 851–865. <https://doi.org/10.1145/3477132.3483557>
- [33] Jeffrey Dean and Luiz André Barroso. 2013. The Tail at Scale. *Commun. ACM* 56 (2013), 74–80. <http://cacm.acm.org/magazines/2013/2/160173-the-tail-at-scale/fulltext>
- [34] Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Vosshall, and Werner Vogels. 2007. Dynamo: Amazon’s Highly Available Key-Value Store. In *Proceedings of Twenty-First ACM SIGOPS Symposium on Operating Systems Principles (Stevenson, Washington, USA) (SOSP '07)*. Association for Computing Machinery, New York, NY, USA, 205–220. <https://doi.org/10.1145/1294261.1294281>
- [35] Dave Dice. 2017. Malthusian Locks. In *Proceedings of the Twelfth European Conference on Computer Systems (Belgrade, Serbia) (EuroSys '17)*. Association for Computing Machinery, New York, NY, USA, 314–327. <https://doi.org/10.1145/3064176.3064203>
- [36] Dave Dice and Alex Kogan. 2019. Compact NUMA-Aware Locks. In *Proceedings of the Fourteenth EuroSys Conference 2019 (Dresden, Germany) (EuroSys '19)*. Association for Computing Machinery, New York, NY, USA, Article 12, 15 pages. <https://doi.org/10.1145/3302424.3303984>
- [37] Dave Dice, Virendra J. Marathe, and Nir Shavit. 2011. Flat-Combining NUMA Locks. In *Proceedings of the Twenty-Third Annual ACM Symposium on Parallelism in Algorithms and Architectures (San Jose, California, USA) (SPAA '11)*. Association for Computing Machinery, New York, NY, USA, 65–74. <https://doi.org/10.1145/1989493.1989502>
- [38] David Dice, Virendra J. Marathe, and Nir Shavit. 2012. Lock Cohorting: A General Technique for Designing NUMA Locks. *SIGPLAN Not.* 47, 8 (Feb. 2012), 247–256. <https://doi.org/10.1145/2370036.2145848>
- [39] Babak Falsafi, Rachid Guerraoui, Javier Picorel, and Vasileios Trigonakis. 2016. Unlocking Energy. In *2016 USENIX Annual Technical Conference, USENIX ATC 2016, Denver, CO, USA, June 22-24, 2016*, Ajay Gulati and Hakim Weatherspoon (Eds.). USENIX Association, 393–406. <https://www.usenix.org/conference/atc16/technical-sessions/presentation/falsafi>
- [40] Songchun Fan and Benjamin C. Lee. 2016. Evaluating asymmetric multiprocessing for mobile applications. In *2016 IEEE International Symposium on Performance Analysis of Systems and Software, ISPASS 2016, Uppsala, Sweden, April 17-19, 2016*. IEEE Computer Society, 235–244. <https://doi.org/10.1109/ISPASS.2016.7482098>
- [41] Panagiota Fatourou and Nikolaos D. Kallimanis. 2011. A Highly-Efficient Wait-Free Universal Construction. In *Proceedings of the Twenty-Third Annual ACM Symposium on Parallelism in Algorithms and Architectures (San Jose, California, USA) (SPAA '11)*. Association for Computing Machinery, New York, NY, USA, 325–334. <https://doi.org/10.1145/1989493.1989549>

- [42] Panagiota Fatourou and Nikolaos D. Kallimanis. 2012. Revisiting the Combining Synchronization Technique. *SIGPLAN Not.* 47, 8 (Feb. 2012), 257–266. <https://doi.org/10.1145/2370036.2145849>
- [43] Jinyu Gu, Qianqian Yu, Xiayang Wang, Zhaoguo Wang, Binyu Zang, Haibing Guan, and Haibo Chen. 2019. Pisces: A Scalable and Efficient Persistent Transactional Memory. In *2019 USENIX Annual Technical Conference, USENIX ATC 2019, Renton, WA, USA, July 10–12, 2019*, Dahlia Malkhi and Dan Tsafir (Eds.). USENIX Association, 913–928. <https://www.usenix.org/conference/atc19/presentation/gu>
- [44] Rachid Guerraoui, Hugo Guiroux, Renaud Lachaize, Vivien Quéma, and Vasileios Trigonakis. 2019. Lock–Unlock: Is That All? A Pragmatic Analysis of Locking in Software Systems. *ACM Trans. Comput. Syst.* 36, 1, Article 1 (March 2019), 149 pages. <https://doi.org/10.1145/3301501>
- [45] Phuong Hoai Ha, Marina Papatrantaflou, and Philippos Tsigas. 2007. Efficient self-tuning spin-locks using competitive analysis. *J. Syst. Softw.* 80, 7 (2007), 1077–1090. <https://doi.org/10.1016/j.jss.2006.11.016>
- [46] Mingzhe Hao, Huaicheng Li, Michael Hao Tong, Chrisma Pakha, Riza O. Suminto, Cesar A. Stuardo, Andrew A. Chien, and Haryadi S. Gunawi. 2017. MittOS: Supporting Millisecond Tail Tolerance with Fast Rejecting SLO-Aware OS Interface. In *Proceedings of the 26th Symposium on Operating Systems Principles* (Shanghai, China) (*SOSP '17*). Association for Computing Machinery, New York, NY, USA, 168–183. <https://doi.org/10.1145/3132747.3132774>
- [47] Danny Hendler, Itai Inceze, Nir Shavit, and Moran Tzafrir. 2010. Flat Combining and the Synchronization-Parallelism Tradeoff. In *Proceedings of the Twenty-Second Annual ACM Symposium on Parallelism in Algorithms and Architectures* (Thira, Santorini, Greece) (*SPAA '10*). Association for Computing Machinery, New York, NY, USA, 355–364. <https://doi.org/10.1145/1810479.1810540>
- [48] Brian Jeff. 2013. big. LITTLE technology moves towards fully heterogeneous global task scheduling. *ARM white paper* (2013).
- [49] Anna R. Karlin, Kai Li, Mark S. Manasse, and Susan S. Owicki. 1991. Empirical Studies of Competitive Spinning for a Shared-Memory Multiprocessor. In *Proceedings of the Thirteenth ACM Symposium on Operating System Principles, SOSP 1991, Asilomar Conference Center, Pacific Grove, California, USA, October 13–16, 1991*, Henry M. Levy (Ed.). ACM, 41–55. <https://doi.org/10.1145/121132.286599>
- [50] Sanidhya Kashyap, Irina Calciu, Xiaohe Cheng, Changwoo Min, and Taesoo Kim. 2019. Scalable and Practical Locking with Shuffling. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles* (Huntsville, Ontario, Canada) (*SOSP '19*). Association for Computing Machinery, 586–599. <https://doi.org/10.1145/3341301.3359629>
- [51] Sanidhya Kashyap, Changwoo Min, and Taesoo Kim. 2017. Scalable NUMA-aware Blocking Synchronization Primitives. In *2017 USENIX Annual Technical Conference, USENIX ATC 2017, Santa Clara, CA, USA, July 12–14, 2017*, Dilma Da Silva and Bryan Ford (Eds.). USENIX Association, 603–615. <https://www.usenix.org/conference/atc17/technical-sessions/presentation/kashyap>
- [52] David Koufaty, Dheeraj Reddy, and Scott Hahn. 2010. Bias Scheduling in Heterogeneous Multi-Core Architectures. In *Proceedings of the 5th European Conference on Computer Systems* (Paris, France) (*EuroSys '10*). Association for Computing Machinery, New York, NY, USA, 125–138. <https://doi.org/10.1145/1755913.1755928>
- [53] Rakesh Kumar, Keith I Farkas, Norman P Jouppi, Parthasarathy Ranganathan, and Dean M Tullsen. 2003. Single-ISA heterogeneous multi-core architectures: The potential for processor power reduction. In *Proceedings. 36th Annual IEEE/ACM International Symposium on Microarchitecture, 2003. MICRO-36. IEEE*, 81–92.
- [54] Rakesh Kumar, Dean M Tullsen, Parthasarathy Ranganathan, Norman P Jouppi, and Keith I Farkas. 2004. Single-ISA heterogeneous multi-core architectures for multithreaded workload performance. In *Proceedings. 31st Annual International Symposium on Computer Architecture, 2004. IEEE*, 64–75.
- [55] Nian Liu, Binyu Zang, and Haibo Chen. 2020. No Barrier in the Road: A Comprehensive Study and Optimization of ARM Barriers. In *Proceedings of the 25th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (San Diego, California) (*PPoPP '20*). Association for Computing Machinery, New York, NY, USA, 348–361. <https://doi.org/10.1145/3332466.3374535>
- [56] Ran Liu, Heng Zhang, and Haibo Chen. 2014. Scalable Read-mostly Synchronization Using Passive Reader-Writer Locks. In *2014 USENIX Annual Technical Conference, USENIX ATC '14, Philadelphia, PA, USA, June 19–20, 2014*, Garth Gibson and Nikolai Zeldovich (Eds.). USENIX Association, 219–230. <https://www.usenix.org/conference/atc14/technical-sessions/presentation/liu>
- [57] David Lo, Liqun Cheng, Rama Govindaraju, Luiz André Barroso, and Christos Kozyrakis. 2014. Towards energy proportionality for large-scale latency-critical workloads. In *ACM/IEEE 41st International Symposium on Computer Architecture, ISCA 2014, Minneapolis, MN, USA, June 14–18, 2014*. IEEE Computer Society, 301–312. <https://doi.org/10.1109/ISCA.2014.6853237>
- [58] Jean-Pierre Lozi, Florian David, Gaël Thomas, Julia L. Lawall, and Gilles Muller. 2012. Remote Core Locking: Migrating Critical-Section Execution to Improve the Performance of Multithreaded Applications. In *2012 USENIX Annual Technical Conference, Boston, MA, USA, June 13–15, 2012*, Gernot Heiser and Wilson C. Hsieh (Eds.). USENIX Association, 65–76. <https://www.usenix.org/conference/atc12/technical-sessions/presentation/lozi>
- [59] Victor Luchangco, Daniel Nussbaum, and Nir Shavit. 2006. A Hierarchical CLH Queue Lock. In *Euro-Par 2006, Parallel Processing, 12th International Euro-Par Conference, Dresden, Germany, August 28 - September 1, 2006, Proceedings (Lecture Notes in Computer Science, Vol. 4128)*, Wolfgang E. Nagel, Wolfgang V. Walter, and Wolfgang Lehner (Eds.). Springer, 801–810. [https://doi.org/10.1007/11823285\\_84](https://doi.org/10.1007/11823285_84)
- [60] Romolo Marotta, Davide Tiriticco, Pierangelo di Sanzo, Alessandro Pellegrini, Bruno Ciciani, and Francesco Quaglia. 2020. Mutable locks: Combining the best of spin and sleep locks. *Concurr. Comput. Pract. Exp.* 32, 22 (2020). <https://doi.org/10.1002/cpe.5858>
- [61] John M. Mellor-Crummey and Michael L. Scott. 1991. Algorithms for Scalable Synchronization on Shared-Memory Multiprocessors. *ACM Trans. Comput. Syst.* 9, 1 (Feb. 1991), 21–65. <https://doi.org/10.1145/103727.103729>
- [62] Yoshihiro Oyama, Kenjiro Taura, and Akinori Yonezawa. 1999. Executing parallel programs with synchronization bottlenecks efficiently. In *Proceedings of the International Workshop on Parallel and Distributed Computing for Symbolic and Irregular Applications*, Vol. 16. Citeseer, 95.
- [63] Zoran Radovic and Erik Hagersten. 2003. Hierarchical Backoff Locks for Nonuniform Communication Architectures. In *Proceedings of the Ninth International Symposium on High-Performance Computer Architecture (HPCA'03), Anaheim, California, USA, February 8–12, 2003*. IEEE Computer Society, 241–252. <https://doi.org/10.1109/HPCA.2003.1183542>
- [64] Sepideh Roghanchi, Jakob Eriksson, and Nilanjana Basu. 2017. Ffwd: Delegation is (Much) Faster than You Think. In *Proceedings of the 26th Symposium on Operating Systems Principles* (Shanghai, China) (*SOSP '17*). Association for Computing Machinery, New York, NY, USA, 342–358. <https://doi.org/10.1145/3132747.3132771>
- [65] Helgi Sigurbjarnarson, James Bornholt, Emina Torlak, and Xi Wang. 2016. Push-Button Verification of File Systems via Crash Refinement. In *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation* (Savannah, GA, USA) (*OSDI'16*). USENIX Association, USA, 1–16.
- [66] Kenzo Van Craeynest, Aamer Jaleel, Lieven Eeckhout, Paolo Narvaez, and Joel Emer. 2012. Scheduling heterogeneous multi-cores through performance impact estimation (PIE). In *2012 39th Annual International Symposium on Computer Architecture (ISCA)*. 213–224. <https://doi.org/>

- 10.1109/ISCA.2012.6237019
- [67] Jons-Tobias Wamhoff, Stephan Diestelhorst, Christof Fetzer, Patrick Marlier, Pascal Felber, and Dave Dice. 2014. The TURBO Diaries: Application-controlled Frequency Scaling Explained. In *2014 USENIX Annual Technical Conference (USENIX ATC 14)*. USENIX Association, Philadelphia, PA, 193–204. <https://www.usenix.org/conference/atc14/technical-sessions/presentation/wamhoff>
- [68] Andrew Wang, Shivaram Venkataraman, Sara Alspaugh, Randy H. Katz, and Ion Stoica. 2012. Cake: enabling high-level SLOs on shared storage systems. In *ACM Symposium on Cloud Computing, SOCC '12, San Jose, CA, USA, October 14–17, 2012*, Michael J. Carey and Steven Hand (Eds.). ACM, 14. <https://doi.org/10.1145/2391229.2391243>
- [69] Xi Yang, Stephen M. Blackburn, and Kathryn S. McKinley. 2016. Elfen Scheduling: Fine-Grain Principled Borrowing from Latency-Critical Workloads Using Simultaneous Multithreading. In *2016 USENIX Annual Technical Conference, USENIX ATC 2016, Denver, CO, USA, June 22–24, 2016*, Ajay Gulati and Hakim Weatherspoon (Eds.). USENIX Association, 309–322. <https://www.usenix.org/conference/atc16/technical-sessions/presentation/yang>
- [70] Kisoo Yu, Donghee Han, Changhwan Youn, Seungkon Hwang, and Jaechul Lee. 2013. Power-aware task scheduling for big. LITTLE mobile processor. In *2013 International SoC Design Conference (ISOC)*. IEEE, 208–212.
- [71] Mingzhe Zhang, Haibo Chen, Luwei Cheng, Francis C. M. Lau, and Cho-Li Wang. 2017. Scalable Adaptive NUMA-Aware Lock. *IEEE Trans. Parallel Distributed Syst.* 28, 6 (2017), 1754–1769. <https://doi.org/10.1109/TPDS.2016.2630695>
- [72] Timothy Zhu, Michael A. Kozuch, and Mor Harchol-Balter. 2017. WorkloadCompactor: reducing datacenter cost while providing tail latency SLO guarantees. In *Proceedings of the 2017 Symposium on Cloud Computing, SoCC 2017, Santa Clara, CA, USA, September 24–27, 2017*. ACM, 598–610. <https://doi.org/10.1145/3127479.3132245>