



# Improving Scalability of Database Systems by Reshaping User Parallel I/O

Ning Li  
Dept. of CSE, University of Texas at  
Arlington  
Arlington, Texas, USA  
ning.li@uta.edu

Hong Jiang  
Dept. of CSE, University of Texas at  
Arlington  
Arlington, Texas, USA  
hong.jiang@uta.edu

Hao Che  
Dept. of CSE, University of Texas at  
Arlington  
Arlington, Texas, USA  
hche@cse.uta.edu

Zhijun Wang  
Dept. of CSE, University of Texas at  
Arlington  
Arlington, Texas, USA  
zhijun.wang@uta.edu

Minh Q. Nguyen  
Faculty of IT, Ho Chi Minh City  
University of Transport  
Ho Chi Minh City, Vietnam  
nqminh@live.com

## Abstract

Modern database systems suffer from compromised throughput, persistent unfair I/O processing and unpredictable, high latency variability of user requests as a result of mismatches between highly scaled user parallel I/O and the I/O capacity afforded by the database and its underlying storage I/O stack. To address this problem, we introduce an efficient user-centric QoS-aware scheduling shim, called AppleS, for user-level fine-grained I/O regulation that delivers the right amount and pattern of user parallel I/O requests to the database system and supports user SLOs with high-level performance isolation and reduced I/O resource contention. It is designed to enable database systems to proactively regulate user request behaviors based on runtime conditions to reshape user access pattern to hide excessive user parallelism from the I/O stack that has a limited concurrent processing capability. This helps achieve scalable throughput for multi-user workloads in a fair and stable manner. AppleS is implemented as a user-space shim for transparent user-differentiated I/O scheduling, making it highly flexible and portable. Our extensive evaluation, run on real databases (MySQL and MongoDB), demonstrates that, by incorporating AppleS in the existing database systems, our solution can not only improve the throughput (up to 39.2%) in a fairer (3.2× to 40.6× fairness improvement) and more stable (up to 2× lower latency variability) manner, but also support user SLOs with less I/O provisioning.

**CCS Concepts:** • Information systems → Autonomous database administration.

**Keywords:** database system, scalability, parallel I/O, QoS

## ACM Reference Format:

Ning Li, Hong Jiang, Hao Che, Zhijun Wang, and Minh Q. Nguyen. 2022. Improving Scalability of Database Systems by Reshaping User Parallel I/O. In *Seventeenth European Conference on Computer Systems (EuroSys '22)*, April 5–8, 2022, RENNES, France. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3492321.3519570>

## 1 Introduction

Although significant progress has been made in improving the transaction throughput scalability for database systems running on multi-core architectures [35, 44, 46, 57, 58, 61, 65], including relational database management systems (RDBMS) for Online Transaction Processing (OLTP), e.g., MySQL [23], Oracle [25] and DB2 [15], and key-value stores, e.g., MongoDB [21], Dynamo [34] and Bigtable [30], modern databases still fail to effectively exploit the available parallelism of concurrent transactions for high-contention multi-user workloads in a fair and stable manner. This makes it challenging for a “database-as-a-service” [54] in a modern datacenter to meet the Service-Level Objectives (SLOs) [29, 50, 55].

Today’s database systems, including relational databases [15, 23, 25] and key-value stores [21, 30, 34], highly rely on buffer cache, lock mechanisms and the storage I/O stack to handle concurrent requests<sup>1</sup>. There user requests that are often arbitrarily delivered through simultaneous user TCP connections may lead to a *user parallelism* beyond the I/O capacity afforded by the database system. Excessive user parallelism can severely compromise database’s I/O performance with *inefficient I/O processing*, *persistent user-level I/O unfairness*, and *unstable request latency*. To demonstrate this, we present, in Fig. 1, the I/O performance of a representative relational database, MySQL 8.0.15 [23], under the baseline,

<sup>1</sup>Request and query are interchangeable in this paper.



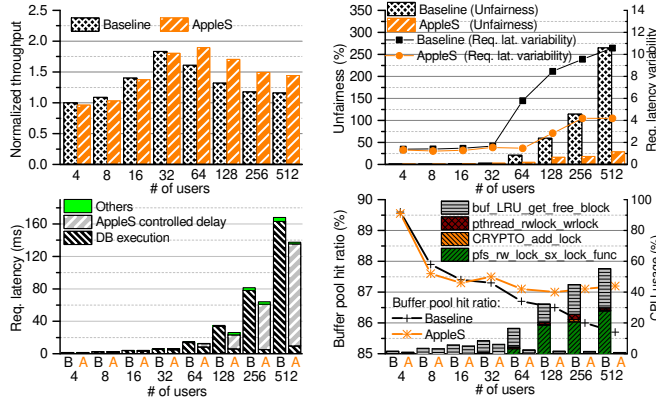
This work is licensed under a Creative Commons Attribution International 4.0 License.

*EuroSys '22, April 5–8, 2022, RENNES, France*

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9162-7/22/04.

<https://doi.org/10.1145/3492321.3519570>




**Figure 1.** Baseline (“B”) results show that request contentions on buffer cache, locks and I/O stack are primary culprits for modern database systems’ inability to provide sufficient transaction throughput, fair I/O processing, low latency, and low latency variability of user requests at high user parallelism. The AppleS (“A”) results show that by temporarily holding back excessive requests outside the database (“AppleS controlled delay”) and controlling the amount and pattern of requests entering the database, such internal contentions can be substantially curbed (“DB execution”), leading to increased throughput, decreased user unfairness, latency and latency variability. The *unfairness* metric measures the maximum deviation in individual user throughput from the average, defined in Section 3.1, which is, the smaller the better.

the latest default I/O scheduler (i.e., mq-deadline [22]) by scaling users. As one can see, the transaction throughput peaks at a specific user parallelism (i.e., around 32 for B, top-left part) when both user I/O unfairness and request latency variability start to deteriorate precipitously (Baseline, top-right part) along with the throughput decline. While the peaking of throughput with increased user parallelism is inevitable and expected at some point, the sudden jump in user unfairness and latency variability is puzzling and begs for more explanations and closer examinations. Given the aforementioned importance of buffer cache, locks and I/O stack in handling user requests, we suspect that request contentions on them at high user-parallelism should be a main culprit, which is confirmed by the baseline results given in the bottom two parts of Fig. 1. That is, starting at user parallelism of 32, the overall request latency and DB execution time rise rapidly (bottom-left part) while the buffer cache hit ratio drops and lock related processing increases notably (black curve, bottom-right part). In other words, the excessive user parallelism significantly lengthens database’s processing time, which is largely spent on the code paths related to locks and buffer cache required to control request concurrency (B bars, bottom-right part). This, along with unbalanced buffer caching across users and queuing effects along the storage I/O stack, directly leads to unstable and uneven user request latency, resulting in persistent user-level I/O unfairness (see Section 2). Our intuitive solution idea,

based on these observations and insight, is to hold back excessive user requests temporarily and control the amount and pattern of requests allowed to enter the database system so that the request contentions on buffer cache, locks and I/O stack are prevented to allow for a “minimally congested” database “internal pipe”. This high-throughput pipe in turn enables all the requests including those held-back requests to get through the database system faster (“AppleS controlled delay”) than if they were not held back (“DB execution”), while the substantially reduced internal request contentions lead to much improved user fairness and latency variability, as evidenced by the “A” results in Fig. 1.

Based on this basic idea, our proposed system in this paper, an *Application level Scheduler* (AppleS), aims to hide excessive user parallelism by strategically holding back the requests surpassing the concurrency afforded by the database system, resulting in lower and stabler request latency, higher throughput, and improved user-level I/O fairness under the given resources.

Unlike the existing solutions, including kernel-level [2, 12, 16, 20, 22, 66] and database-specific approaches [42, 45, 48, 51, 52, 56, 60], that passively react to one or more bottlenecks resulting from highly scaled user parallelism by improving some internals of a database or instrumenting OS-kernel modules, AppleS is designed to nip the aforementioned problem in the bud by proactively reshaping user request streams in such a way that it prevents excessive I/O workload (beyond the I/O capacity afforded by the database and its underlying storage I/O stack) from reaching and overloading the database system. In other words, AppleS trades controllable and predictable delays of excessive user requests waiting outside the database system for an uncongested and resource-efficient runtime database system without which their I/O delays inside the database system can become unpredictably long and variable, harmful to user I/O fairness and efficiency. Specifically, to capture the relationship between user behaviors and a database system’s performance, we regulate the user parallel I/O perceived by and delivered to the database system in two dimensions, i.e., *delivered user parallelism* and *scheduling-round I/O quota*. The former is the number of concurrent users that are seen by the database system actively contending for I/O resources, while the latter determines the number of consecutive requests issued in a scheduling round for each user according to its priority. The intuition is that, a value of the former matching the inherent concurrency of the database system allows the most effective usage of the system resources and hence, optimize throughput, whereas, a value of the latter based on user priority would help achieve this throughput without sacrificing user-level fairness by exploiting per-user sequentiality and locality (e.g., for reduced cache misses). By optimizing the values of delivered user parallelism and scheduling-round I/O quota at runtime, AppleS is able to effectively avoid or alleviate the performance bottlenecks



**Figure 2.** Each data point represents the cumulative wait time experienced by a given user, which issues the TPC-C workload to access a MySQL 8.0.15 instance. The unbalanced lock waits distributed across users can be aggravated by highly scaled user parallelism.

due to superfluous user parallel I/O. This enables AppleS to support user QoS with effective user-level isolation and reduced I/O provisioning. In addition, AppleS’s non-intrusive implementation, without the need to modify databases and OS kernels, makes it compatible with and portable to different types/versions of databases and different versions of OS kernels.

We implement an AppleS prototype and evaluate its effectiveness under four versions of the Linux kernel and compare it against nine state-of-the-art or widely accepted I/O schedulers and a Linux resource management tool (i.e., Cgroup(v2) [3]) with extensive experiments driven by two types of databases, including a RDBMS, MySQL [23] (version 8.0.15 and 8.0.23), for OLTP evaluation, and a key-value store, MongoDB [25] (version 3.6.0 and 4.4.3), for cloud service workloads [32]. Our prototype evaluation results demonstrate that AppleS is able to provide high transaction throughput, significantly improve on user-level I/O fairness (up to 40.6 $\times$ ) and maintain a low request latency and latency variability when the number of concurrent users scales. More importantly, AppleS can integrate other QoS guarantee algorithms (e.g., PSLO [49]) to enforce customized user-level SLOs with reduced I/O provisioning (saving up to 53.1% transaction throughput) by supporting fairer and more stable I/O resource allocation to concurrent users. The source code of AppleS is now available at [https://github.com/NingBellWind/AppleS\\_Artifact/](https://github.com/NingBellWind/AppleS_Artifact/).

## 2 Background and Motivation

This section first presents the necessary background to set the stage for the analysis and empirical observations later in the section that motivate the AppleS study.

### 2.1 Background and Case Studies

Highly scaled *user parallelism*<sup>2</sup> can cause scattered performance bottlenecks (e.g., the buffer cache, lock mechanisms of the database, the kernel-level block layer, etc.) along the whole *I/O stack* that consists of the database and the underlying storage I/O stack. To examine these performance bottlenecks aggravated by high user parallelism, we first discuss their impact on user-level I/O performance isolation, and then study two specific cases of databases, a relational

database, MySQL 8.0.15 [23], and a key-value store, MongoDB 3.6.0 [21], to quantify the impact of different levels of user parallelism on transaction throughput, user-level I/O fairness and request-latency variability. This helps us reveal the root causes behind these performance bottlenecks.

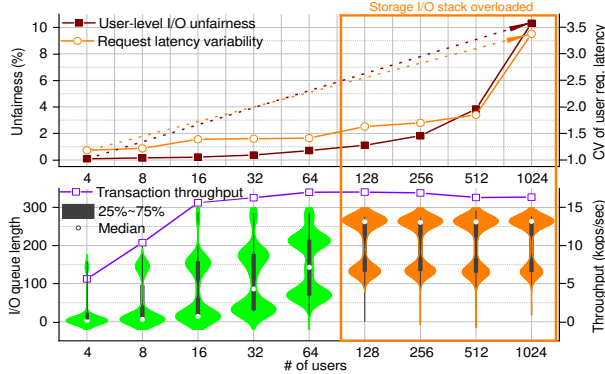
**User I/O Performance Isolation for Databases:** As a key indicator of the lack of user I/O performance isolation, user-level I/O unfairness (defined in Section 3.1) can be attributed to many factors and chief among them are cache conflicts and lock contentions in the database and storage I/O queuing effects. Cache conflicts can induce cache residency imbalance [36] where faster running users are rewarded with more buffer cache space, leading to their increased throughput allocations but at the expense of slower-running users. This vicious cycle causes unacceptable *persistent user-level I/O unfairness*. Moreover, existing transaction-based lock management mechanisms [46, 53, 65] cannot effectively track and prevent likely user-level unfairness resulting from intensified contentions on database locks (e.g., by invoking `pfs_rw_lock_sx_lock_func` in MySQL [23]), as evidenced in Fig. 2. Similarly, other solutions focusing on databases’ caching mechanisms [42, 45, 52, 56] also cannot effectively track and regulate user-level I/O traffic, falling short of supporting user-level fair and stable I/O sharing. In addition, storage I/O queuing effects can further magnify user I/O unfairness [43]. These three factors often work together to exacerbate the problem of persistent user I/O unfairness under high user parallelism. However, for different types of databases, they can play different roles, as elaborated next.

**Bottlenecks of MySQL:** We investigate the performance of MySQL 8.0.15 to understand its bottlenecks. With the user parallelism continuously increasing beyond the I/O concurrency afforded by the database, lock contentions and buffer cache conflicts become increasingly more frequent and serious (Baseline, Fig. 1), leading to inefficient usage of resources throughout the I/O stack and hence, decreased transaction throughput and increased user-level I/O unfairness and request-latency variability (Baseline, Fig. 1).

The intensified lock contentions and cache conflicts by the scaling of concurrent users cause the transaction throughput to decrease when the number of users surpasses a threshold (i.e., 32), as illustrated in Fig. 1. In the meantime, I/O resources allocated to different users are increasingly unbalanced as user parallelism increases, as indicated by the unfairness measure and request-latency variability in Fig. 1.

<sup>2</sup>It refers to the number of concurrent I/O-intensive users contending for I/O resources.





**Figure 3.** Challenges posed by excessive user parallelism for MongoDB. The violin plots show the distributions of I/O queue length observed in the I/O stack as a function of user parallelism. The median queue length and the interquartile range (from 25% to 75%) for a given user parallelism are represented by the white dot and the thick black vertical bar respectively.

It is noted that `pfs_rw_lock_sx_lock_func` consumes more CPU cycles when the number of concurrent users goes beyond a threshold (i.e., 32), which implies more expected processing delay due to row-level shard and exclusive locks [19] with more concurrent users. Similarly, a high CPU usage is observed on `buf_LRU_get_free_block` when the number of users exceeds 32. This means that the buffer pool frequently evicts LRU pages and flushes them to the disk to accommodate new pages, aggravating cache conflicts, resulting in reduced buffer pool hit ratio.

In contrast, by trading a controlled delay at the socket buffers, which can increase with the number of concurrent users scaling up, for a reduced latency inside the database-I/O-stack blackbox, AppleS achieves higher throughput and overall lower latency than uncontrolled situations, as illustrated in Fig. 1.

**Bottlenecks of MongoDB:** I/O queuing effects along the storage I/O stack can become a major bottleneck under high user parallelism for MongoDB [21]. We verify this effect by exponentially increasing the number of I/O-intensive concurrent users imposed on MongoDB 3.6.0 with the YCSB [32] workloads (configurations are given in Section 5) under a state-of-the-art I/O scheduling framework (Split-level [66]). As shown in Fig. 3, a larger number of users can lead to a higher queue-length variability in the storage I/O stack and thus an unstable I/O service. Even with the Actually Fair Queuing (AFQ) scheduler [66], user I/O unfairness and request-latency variability (by measuring the coefficient of variation (CV) [13] averaged across all users) increase by 100× and 3× respectively from 4 users to 1024 users.

**Conclusions:** As the storage backend supporting cloud services for multi-tenancy [37, 54] and large-scale industrial services [17, 18], database systems are often faced with high user parallelism. However, for both MySQL [23] and MongoDB [21], highly scaled parallel I/O can significantly weaken user I/O performance isolation and/or degrade I/O

efficiency. The key is to fully exploit the potentials of database systems to service more concurrent users by minimizing the side effects, i.e., lock contentions, cache conflicts and/or queuing effects, aggravated by high user parallelism.

## 2.2 Related Work

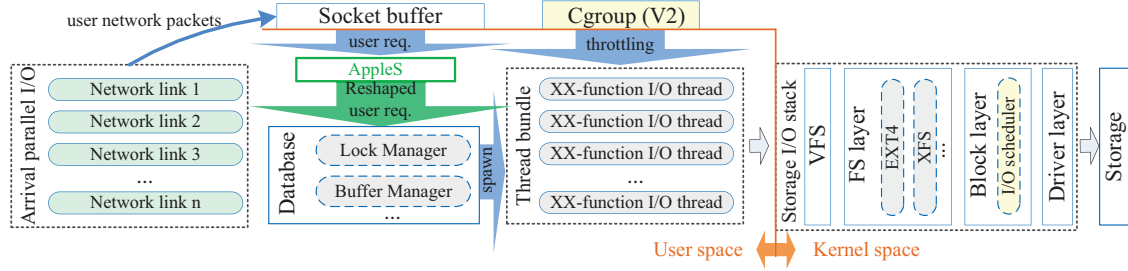
In this section, we assess the existing solutions addressing the challenges due to excessive user parallelism for database systems, which further motivates our work.

**Kernel-Level I/O Schedulers:** The I/O schedulers [12, 16, 20, 22, 66] perform different levels of support for *process/thread-level* fairness and bursty I/O control in a passive manner. However, these schedulers cannot identify the user visiting the application that produced the I/O request and thus cannot effectively track and regulate user-level I/O traffic for databases.

**Resource Management Tools:** Existing resource management tools, e.g., Cgroups [2], can provide great flexibility in process/thread-level resource allocation. However, like I/O schedulers, Cgroups, as an OS kernel feature, cannot effectively track and regulate user-level I/O traffic that happens on top of the database.

**Database-Specific Solutions:** Existing efforts on the transaction scalability of OLTP mainly focus on efficiency problems stemming from database internals such as contentions on concurrency control and the buffer cache. For example, many solutions improve the caching or buffering mechanism by optimizing replacement algorithms [42, 45, 52, 56] or buffer cache management [48, 51, 60] for high transaction throughput or low latency. Other solutions [39, 44, 46, 57, 58, 61, 65] aim to enhance the capability of concurrency control for achieving optimized transaction scalability under high user parallelism. However, these solutions are not designed to support user-level fair and stable I/O sharing. Other database-specific solutions (e.g., SILK [28], Rein [59] and DAST [31], etc.) essentially incorporate improved components into a specific database or an evaluation framework (e.g., Janus codebase [4]) for performance improvement, still incapable of proactively alleviating the performance bottlenecks in the storage I/O stack and regulating user parallel I/O behaviors before they enter the database, which is proven to be the key to minimizing the side effects caused by excessive user parallelism (See Fig. 1). More importantly, because databases and OS kernels are frequently updated, it can entail a substantial amount of work (e.g., coding, testing and integration) for a database-specific solution to migrate across different versions of databases and/or OS kernels. This largely prevents database-specific solutions from offering the necessary portability and compatibility to different versions/types of databases and OS kernels.

In contrast, AppleS can effectively alleviate the aforementioned performance bottlenecks without instrumenting databases or OS kernels. Meanwhile, AppleS is orthogonal and complementary to the database internal solutions and



**Figure 4.** Overview of AppleS. As a user parallel I/O regulator, AppleS reshapes the access pattern of user requests before they enter the database, which makes AppleS potentially orthogonal to and compatible with database-specific and kernel-level solutions.

hence, can work seamless with different types/versions of databases for further performance enhancement, as demonstrated in Section 5.

### 3 AppleS Design

AppleS works as a user I/O regulator on top of the *I/O stack*, that consists of the database and its underlying *storage I/O stack*, as illustrated in Figure 4. A database running instance (e.g., for MongoDB [21] and MySQL [23]) typically adopts a multi-thread model to handle user requests wrapped as the corresponding network packets, which are received from users' network links and temporarily stored in the OS socket buffer. The storage I/O stack, including virtual file system, file system layer, block layer and driver layer, etc., leverages multi-threaded I/O requests to fairly and efficiently share the storage backend. *AppleS works as a user-space scheduling shim treating the whole I/O stack as a black box.*

#### 3.1 Design Goals and Principles

"Database-as-a-service" [54] in a modern datacenter demands high *throughput* (i.e.,  $\lambda$ ) under effective *performance isolation*, e.g., in terms of *user-level I/O unfairness* (i.e.,  $U_f$ ) and/or *request latency variability* (i.e.,  $V_l$ ), to serve concurrent users in a fair, stable and efficient fashion.  $U_f$  is a dimensionless metric to assess relative I/O fairness between any two backlogged users ( $u_i$  and  $u_j$ ,  $1 \leq i, j \leq n$ ) during any time interval  $[t_1, t_2]$ .  $U_f = \frac{\max_{i=1}^n |N_i(t_1, t_2) / \phi_i - H|}{H}$ , where  $H = (\sum_{i=1}^n \frac{N_i(t_1, t_2)}{\phi_i}) / n$ ,  $N_i(t_1, t_2)$  is the number of scheduled requests for user  $u_i$ ,  $1 \leq i \leq n$ , during  $[t_1, t_2]$ . Intuitively and informally,  $U_f$  is the maximum deviation in individual user throughput from the average. To make the metric fit for users' transactions with distinct loads,  $U_f$  is measured at the request/query level.  $V_l$  is defined as the coefficient of variation (CV) [13] of request latency averaged across users. Our AppleS targets at maximizing total I/O throughput under strong performance isolation across highly scaled concurrent users. To this end, the AppleS design is guided by the following principles:


**Early intervention:** Highly scaled user parallel I/O can aggravate cache conflicts, lock contentions and the I/O queuing effects, as shown in Fig. 1-3. Rather than individually addressing these bottlenecks, AppleS aims to proactively

reshape user parallel I/O before it causes the spread of bottlenecks along the I/O stack.

**Hiding the excessive I/O:** AppleS aims to hide excessive user parallel I/O beyond the I/O-stack capacity by delivering the right amount of user parallelism to and keeping the appropriate per-user consecutive access pattern for the I/O stack at runtime.

**Portability and compatibility:** With the only system assumption being user-level network-IO syscall interception, AppleS is portable to different versions/types of databases (i.e., MongoDB [21] and MySQL [23]) by simple reconfigurations. In addition, AppleS is designed to be compatible with different OS kernels and orthogonal to existing kernel-level resource management tools (e.g., Cgroups [2]) and I/O schedulers [12, 16, 20, 22, 66].

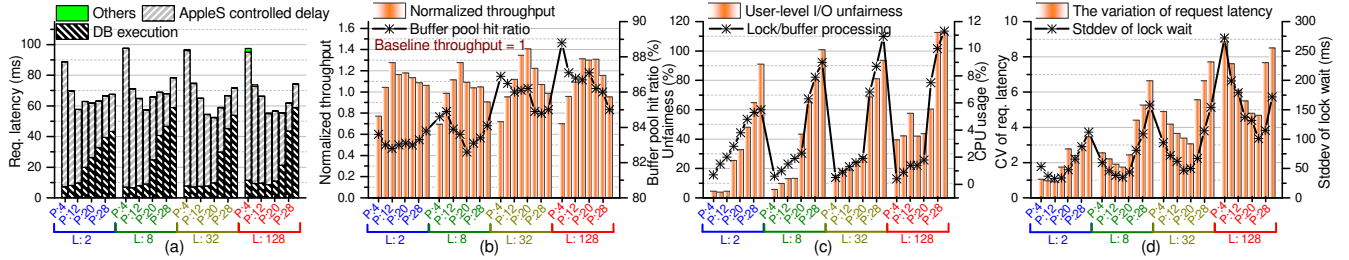
#### 3.2 Early Intervention



**Figure 5.** The life cycle of user I/O requests.

It is desirable but challenging to accommodate excessive user requests beyond the I/O-stack capacity before they enter the I/O stack. Without instrumenting the database and OS kernel as traditional designs and implementations do, AppleS regulates user parallel I/O by intervening the process that transforms the network I/O packets from user network links into the user I/O requests to be handled by the database.

Figure 5 illustrates the life cycle of user I/O requests by showing their 6-step process flow. User I/O requests wrapped in I/O packets are first received by the NIC driver and then stored at a socket buffer in the Linux kernel (step 1). The database continuously issues network-I/O system calls to signal if I/O packets are ready for the database (step 2). If yes, I/O packets will enter the database and be transformed into user requests for business-logic processing (step 3). Then, these user requests will be converted into thread-level I/O requests for the OS kernel and the rest of the I/O stack (step 4). Once the results for these requests are fed back from the Linux kernel (step 5), the database will process them and send these results back to the users (step 6). Intuitively, by



**Figure 6.** The effectiveness of the AppleS scheduling under different combinations of delivered user parallelism  $P$  and scheduling-round I/O quota  $\zeta$ . For the convenience of observing the impact of per-user I/O sequentiality on I/O performance, which is controlled by  $L_i$  for 256 concurrent users,  $1 \leq i \leq 256$ , each user is assigned the same weight  $\phi^0$  and thus the same value of  $L = \zeta * \phi^0$  for  $L_i$ .

strategically intercepting and suspending network-I/O system calls, AppleS can effectively shield the user requests that exceed the I/O-stack processing capacity from overloading any part of the I/O stack.

### 3.3 Hiding Excessive I/O

AppleS reshapes the user parallel I/O delivered to the database system by two control knobs, i.e., *delivered user parallelism*, denoted as  $P$ , and *scheduling-round I/O quota*, denoted as  $\zeta$ , used by a runtime I/O scheduler.

**Scheduling:** AppleS schedules user requests under a weighted round-robin (WRR) policy [27, 47], aiming to coordinate user parallelism and access pattern in a fine-grained fashion. The I/O scheduling during a period of time can be logically divided into multiple *scheduling rounds*. During each round, the user parallelism seen by the I/O stack is limited by  $P$  while the total number of user requests to be issued is bounded by  $\zeta$ . The *user I/O quota* will be assigned according to their individual weights  $\phi_i$  to implement I/O differentiation under the fairness policy, i.e.,  $L_i = \phi_i * \zeta$  for the  $i^{th}$  user,  $u_i$ . Specifically, the per-user quota limits the number of requests a user is allowed to issue in a scheduling round. If some users don't use up their quotas, all others would benefit because their requests will be moved up to take over the unused quota based on the WRR scheduling policy. In doing so, AppleS can coordinate I/O rate allocation across concurrent users to meet user QoS targets.

The core idea behind the AppleS scheduling is to coordinate the user request latency spent outside of the database system by controlling  $P$ , i.e., the AppleS controlled delay, which can considerably reduce the request delay inside the database resulting from buffer cache conflicts or lock contentions that can be intensified by the excessive user parallelism inside the database. In addition, AppleS can also enable per-user I/O sequentiality adjustment for exploiting potential I/O locality (e.g., for reduced cache misses) by increasing  $\zeta$ .

In contrast to existing admission control systems [38, 41, 68, 69], AppleS selectively delays excessive requests from multiple concurrent users to maintain the matching point (i.e.,  $P$  control), under which users are not allowed to issue more requests than their respective quotas in a scheduling

round (i.e.,  $\zeta$  control). Since AppleS maximizes throughput, the overall delays experienced by the held-back requests will be shorter than those by requests under conventional throttling or without throttling that reduces throughput as user-parallelism increases beyond a threshold (as evidenced by Fig. 1). Pending requests are withheld in the socket buffers, consuming minimum resources.

To show the effect of the AppleS scheduling with the two control knobs and provide an intuition for finding their optimal combination later in the paper, we let 256 users, which issue the TPC-C workload [26], concurrently access a MySQL 8.0.15 instance controlled by AppleS under different combinations of  $P$  and  $\zeta$ . To effectively investigate the impact of per-user I/O sequentiality on I/O performance, which is controlled by  $L_i$ , each user has the same weight  $\phi^0$  and thus the same value of  $L = \zeta * \phi^0$  for  $L_i$ . According to our experimental observations, we can lay out the impact of  $P$ - $\zeta$  control on I/O performance as follows:

**$P$  Control:** The  $P$  control is the key to regulating user-level parallel I/O workload to match the I/O concurrency afforded by the database system. It affects the AppleS controlled delay experienced by those excessive requests temporarily held back because they were considered beyond the I/O-stack capacity based on the  $P$  limit. As shown in Fig. 6(a), there is a negative correlation between  $P$  and the AppleS controlled delay. A larger  $P$  commonly corresponds to a shorter AppleS controlled delay, due to the relaxation on user I/O concurrency limit, but likely causes suddenly surged database execution, largely because of a sharp CPU usage increase on lock/buffer-related processing (see in Fig. 6(c)), e.g., by frequently invoking `pfs_rw_lock_sx_lock_func` and `buf_LRU_get_free_block` in MySQL [23] (recall Section 2.1). Thus, the  $P$  control can help seek a *matching point* with the shortest AppleS controlled delay just before database execution time surges, resulting in a maximized throughput (see in Fig. 6(b)). Note that the  $P$  value with the maximized throughput commonly corresponds a low-level variability in request latency and lock waits (see in Fig. 6(d)). This means that running at the matching point will likely contribute to an efficient and stable I/O sharing for concurrent users. In addition, for different  $L$  (and thus  $\zeta$ ), their matching points basically stay close to each other (i.e. the

range of  $P$  between 16 and 20) except the case of the smallest I/O sequentiality<sup>3</sup> (i.e.,  $L = 2$ ) where the matching point is located at the range of  $P$  between 12 and 16. *This indicates a matching point obtained under a specific level of per-user I/O sequentiality can be applied to other levels of I/O sequentiality.*

**$\zeta$  Control:** A larger  $\zeta$  can help exploit per-user I/O sequentiality and locality for a higher buffer cache hit ratio provided that  $P$  does not exceed the matching point (see in Fig. 6(b)). However, too large a  $\zeta$  (e.g., when  $L = 128$  and thus  $\zeta = 128 * 256$ ) can lengthen lock waits and offset the throughput gains from the enhanced cache hit ratio (see in Fig. 6(a)(b)). This is because a higher per-user I/O sequentiality indicates that more requests from the same user can enter the database in the same scheduling round, implying a potentially higher data locality among these requests (e.g., accessing the same data row) and resulting in increased row-lock time. Moreover, too large a  $\zeta$  can also aggravate user I/O unfairness and request latency variability due to aggravated cache residency imbalance [36] (recall Section 2.1) and the intensified variation of lock time, respectively (see in Fig. 6(c)(d)). Finding the right  $\zeta$  can help maximize throughput under the constraints of user I/O fairness and request latency variability. In addition, a smaller  $P$  means lower user-parallelism delivered to the database, leading to a higher hit-ratio (due to fewer conflicts) but also reduced I/O-stack utilization. The matching point guides  $P$  control to shield excessive user requests while fully exploiting I/O resources, the prerequisite for the  $\zeta$  control for a high hit-ratio. It is noted that the  $\zeta$  control has almost negligible impact on database execution and AppleS controlled delay (see in Fig. 6(a)), which makes the  $\zeta$  control nearly orthogonal to the  $P$  control. This is because the  $\zeta$  control, which regulates the per-user I/O interarrival time distribution, has little impact on the user parallelism (across users) delivered to the database system once  $P$  is determined.

As a result, a key  **$P$ - $\zeta$  control principle** is that *the  $P$  control that curbs excessive user parallelism takes precedence to and is a precondition for the effectiveness of the  $\zeta$  control that attempts to increase cache hit ratio by increasing per-user access locality, while the  $\zeta$  control has negligible impact on the delivered user parallelism controlled by  $P$ .* Based on this principle, we can start by optimizing  $P$  control for obtaining the matching point at a specific  $\zeta$  setting and then further optimize  $\zeta$  control to maximize throughput constrained by the requirements of user I/O fairness and/or request latency variability.

### 3.4 Portability and Compatibility

First, AppleS is designed to regulate user parallel I/O on top of the database and the underlying storage I/O stack by intercepting network-I/O system calls, essentially treating the I/O stack as a black box. This separates AppleS's operations from

the underlying database processing as well as kernel-level I/O scheduling and the process/thread-level resource throttling executed by resource management tools (e.g., Cgroup (V2) [2]).

Second, from the perspective of AppleS, the difference among different types of databases only lies in the adoption of different types of network-I/O system calls to handle user requests. Thus, to enable AppleS to work for a specific database, we only need to configure AppleS with the identification codes of the system calls associated with user request processing, while the database's internals can change greatly from version to version. As a blackbox solution, AppleS will automatically coordinate the  $P/\zeta$  controls to optimally adapt to these changes.

Finally, if load balancing middleware is available for the database system, a joint load-balance control in the middleware and the AppleS control in individual backend database servers can be devised to enhance system-level I/O performance. This is because AppleS can support efficient, fair and stable I/O sharing under high user parallelism for each server, thus further improving the system-level capability to service more concurrent I/O intensive users by complementing load balancers.

## 4 AppleS Implementation

The AppleS implementation consists of three main functional modules, namely, the *User-context builder*, *Scheduler* and *Optimizer*. The user-context builder is responsible for establishing the per-user request queues; the scheduler reshapes user access pattern according to  $P/\zeta$  controls. As the brain of AppleS, the optimizer tunes  $P$  and  $\zeta$  through a two-timescale tuning process, slow for  $P$  and fast for  $\zeta$ , to provide user-level I/O sharing in a fair, stable and efficient manner.

### 4.1 User-Context Builder

The user-context builder module builds user-level context by extracting the user-level network connection information hidden behind user requests. The module analyzes the system calls from user requests to extract the peer IP address and port for each user connection. Thus, a user in AppleS can be uniquely identified by its  $\langle IP\text{-}address, port\text{-}number \rangle$  tuple. With the user-specific ID information so obtained, a user-level request queue is built for each distinct user by grouping requests from that user in the FIFO order.

### 4.2 Scheduler

AppleS scheduler, implemented as a user-space module based on a `syscall_intercept` library [5], intercepts and schedules network-I/O system calls issued by the database. The library allows AppleS to hook the Linux system calls in the user space by hot-patching the machine code of the library in the memory of a process, enabling AppleS to intercept and schedule them. When AppleS delays the system calls, the I/O

<sup>3</sup>The case of  $L = 1$  is considered no I/O sequentiality.



threads waiting for their feedback will be blocked until AppleS releases its control. Thus, all the user-level scheduling goals can be realized for a specific database by intercepting user requests destined for the I/O stack. The overhead of the system-call interception is found to be negligibly small, i.e., on the order of 100 nanoseconds, or larger than 1 million interceptions/second. Given the sub-Million IOPS throughput of the existing storage volume for high-performance databases (e.g., 300K IOPS per volume for the Ultra high-performance Oracle cloud infrastructure block volumes reported in June 2021 [7]), we believe that AppleS should be scalable to higher-performance databases or soft-realtime applications [33].

### 4.3 Optimizer

Based on the analysis in Section 3.3, the optimizer workflow can be broadly divided into two parts, namely, optimizations for the control knobs  $P$  and  $\zeta$  respectively. The I/O-stack concurrency capability  $C$  (i.e., the optimal value of  $P$ ) is determined by the resource provision in the I/O stack and thus remains relatively stable since resource configuration changes infrequently. Moreover, the impact of fast changing workloads on user-experienced I/O performance can be largely regulated by AppleS in a stable and efficient manner under the matching point since instantaneously excessive concurrent requests can be effectively held back out of the I/O stack. Thus, the  $P$  optimization runs at a slow timescale (e.g., minute-level timescale). In contrast, the  $\zeta$  optimization, which aims to regulate and exploit per-user request sequentiality, must react quickly to the instantaneous I/O bursts to meet the user-level requirements for I/O fairness or request-latency variability, and therefore operates at a fast timescale (e.g., 100ms-level) provided that each backlogged user can issue at least one request during each scheduling round.

**4.3.1  $P$  Optimization.** The goal is to find the smallest  $P$  where the throughput (i.e.,  $\lambda$ ) peaks, i.e., the matching point,  $C$ , that support high throughput under effective user I/O performance isolation. To keep the measurement complexity low, we resort to a hybrid modeling-and-measurement-based approach. Specifically, by experiment, we first found that the percentage of parallel execution time,  $\alpha(P)$ , also known as parallel fraction, can be easily modeled as an exponentially decreasing function of  $P$  with three parameters. Then with a simple AppleS execution model, we are able to establish  $\lambda(P)$  as an explicit function of  $P$  via  $\alpha(P)$ , with the three parameters underpinned by only three throughput samples, resulting in extremely low measurement complexity of  $O(1)$ . The peak of  $\lambda(P)$  is then identified with high accuracy as the experiments show.

By experiment, as shown in Fig. 7, we found that  $\alpha(P)$  closely follows an exponentially decreasing function with a non-zero floor. In other words, it can be generally modeled as:




Figure 7.  $\alpha$  curve fitting for MongoDB and MySQL.

$$\alpha(P) = A + B * e^{-\gamma * P}, \quad A, B, \gamma > 0 \quad (1)$$

with only three positive parameters,  $A$ ,  $B$  and  $\gamma$ . Fig. 7 shows the curve fitting of  $\alpha(P)$  given by Eq. 1. We observe that both fitting curves closely track real data points (i.e., using throughputs measured at different  $P$  values for MongoDB [21] and MySQL [23] databases, under the YCSB [32] and TPC-C benchmark [26] workloads, respectively) and accurately reflect their trend.


With the empirical expression of  $\alpha(P)$  in Eq. 1, we show in Appendix that with a simplified execution model that approximately captures the AppleS execution process, the throughput  $\lambda(P)$  can be given as,

$$\lambda(P) = \frac{P * \lambda(P_0)}{P_0 + B * P_0 * (e^{-\gamma P_0} - e^{-\gamma P}) * (P - 1)}. \quad (2)$$

To pin down the parameters in Eq. 2, including  $\lambda(P_0)$ ,  $P_0$ ,  $B$ , and  $\gamma$ , we resort to a measurement approach by sampling three throughput data points. The first data point is  $\lambda(P_0)$  at  $P = P_0 (< C)$ . In this case,  $\alpha$  is expected to approach 1. This is because with a delivered user parallelism smaller than the I/O concurrency capacity of the I/O stack, the I/O stack may be underutilized and virtually all issued requests can be processed concurrently. Now, there are two parameters, i.e.,  $B$  and  $\gamma$ , left to be fixed, which calls for two more samples to be measured. Specifically, with two throughput values,  $\lambda(P_1)$  and  $\lambda(P_2)$ , sampled at  $P = P_1$  and  $P = P_2$ , respectively, as input to Eq. 2, we establish two equations which can be jointly solved to uniquely determine  $B$  and  $\gamma$ .  $\lambda(P_2)$  is sampled without constraining  $P$  (i.e.,  $P_2$  = the number of users) while  $\lambda(P_1)$  can be sampled at a medium value. In principle, the  $\lambda(P_1)$  sample can be collected at any  $P_1$  values. Nevertheless, we recommend to sample them (i.e.,  $\lambda(P_1)$  and  $\lambda(P_2)$ ) far apart to reflect the throughput variation with  $P$ . Finally, the matching point,  $P = C$ , is found as the first local maximum of  $\lambda(P)$  given by Eq. 2.

Algorithm 1 further discloses the implementation details of the  $P$  optimization, which mainly consists of three parts, i.e., three-data-point throughput sampling, pinning down the parameters  $B$  and  $\gamma$ , and finding the matching point. As the inputs for the  $P$  optimization, the throughput measures are required to represent the burstiness and load added by not only the database's processing but also background I/O-intensive tasks that can significantly affect database performance especially under high user parallelism, e.g., *fsync()*-induced





**Figure 8.** A comparison between AppleS estimates for throughput based on different samplings of the three throughput data points and the actual measurements.

---

**Algorithm 1:** P Optimization Algorithm

---

**Input:** The delivered user parallelism (i.e.,  $P_0, P_1$ , and  $P_2$ ) for three-data-point sampling  
**Output:** The matching point (i.e.,  $C$ )

```

1 /*Three-data-point throughput sampling*/
2  $\lambda(P_0) \leftarrow \text{throughput\_measure}(P_0)$ 
3  $\lambda(P_1) \leftarrow \text{throughput\_measure}(P_1)$ 
4  $\lambda(P_2) \leftarrow \text{throughput\_measure}(P_2)$ 
5 /*With three measured throughputs,  $\lambda(P_0), \lambda(P_1)$ , and
    $\lambda(P_2)$ , the standard solver based on the Newton-Raphson
   method is used to solve for  $\gamma$  from
    $f(\gamma) = B * P_0 * (e^{-\gamma P_0} - e^{-\gamma P_1}) * (P_1 - 1) - P_1 * \lambda(P_0) / \lambda(P_1)$ 
   (derived from Eq. 2).*/
6  $\gamma \leftarrow \text{Newton\_Raphson}(f(\gamma))$ 
7 /*Determine  $B$  with three data points based on Eq. 2*/
8  $B = \frac{P_1 * \lambda(P_0) / \lambda(P_1) - P_2 * \lambda(P_0) / \lambda(P_2)}{P_0 * ((e^{-\gamma P_0} - e^{-\gamma P_1}) * (P_1 - 1) - (e^{-\gamma P_0} - e^{-\gamma P_2}) * (P_2 - 1))}$ 
9 /*Find the smallest  $P$  where the throughput peaks, i.e., the
   matching point*/
10  $P \leftarrow 1$ 
11  $\text{last\_thpt} \leftarrow 0$ 
12  $\text{curr\_thpt} \leftarrow \lambda(1)$ 
13 while  $\text{curr\_thpt} > \text{last\_thpt}$  do
14    $\text{last\_thpt} \leftarrow \text{curr\_thpt}$ 
15    $P \leftarrow P + 1$ 
16    $\text{curr\_thpt} \leftarrow \lambda(P)$ 
17 return  $P - 1$ 
```

---

data flushing. Databases need to periodically (e.g., 60 seconds by default for MongoDB [9]) flush pending writes to the storage backend (e.g., by the *fsync()* system call), which can cause the "fsync freeze" problem [1, 66] and compromise database transaction throughput and latency performance. Thus, the interval length for each measurement is set at 60 seconds by default. Based on the three-data-point measurements, one can solve  $B$  and  $\gamma$  by the Newton-Raphson method [24]. Specifically, Algorithm 1 first derives a function of  $\gamma$  (i.e.,  $f(\gamma)$ ) based on Eq. 2. And then, it conducts multiple iterations, each of which updates  $\gamma$  with a better approximation (i.e., the x-intercept of tangent line obtained by  $\gamma - f(\gamma)/f'(\gamma)$ ), to obtain a  $\gamma$  with a desired accuracy, which is guaranteed by  $|f(\gamma)| < \mu$  ( $\mu$  is set at  $10^{-12}$  by default) only if the condition for convergence of the Newton-Raphson

method can be met. After that,  $B$  can be solved based on Eq. 2 with  $\lambda(P_0), \lambda(P_1)$  and  $\lambda(P_2)$  as input. Finally, Algorithm 1 can check the throughput obtained from Eq. 2 under increasing levels of  $P$  from 1 until the peak of throughput can be identified, which is the matching point.

We further verify the precision of Eq. 2 on a MySQL 8.0.15 instance. Specifically, we let 128 users, which issue TPC-C workload [26], concurrently access the MySQL instance controlled by AppleS under distinct  $\zeta$  settings. We assess the percentage error [11] between the measured throughput at the estimated matching point derived from Eq. 2 and the throughput obtained under the actual matching point, i.e.,  $\lambda_{\max}$  error. We also evaluate the sensitivity of Eq. 2 by choosing different middle samples,  $\lambda(P_1)$ . As shown in Fig. 8, Eq. 2 can accurately estimate the trend of throughput as a function of  $P$  under different  $\zeta$  settings, which can be confirmed by the average  $\lambda_{\max}$  error (i.e., 3.58%) over all the cases in Fig. 8. More importantly, Eq. 2 is shown to be insensitive to the choice of middle samples as different middle samples, which lead to throughput trends with negligible difference, indicating a satisfactory stability for the output when the input changes (e.g., algorithmic stability [10]). Similarly, we also observe a low  $\lambda_{\max}$  error for MongoDB [21] with the YCSB workloads [32] (detailed configurations in Section 5), e.g., for the maximum number of users (i.e., 1024) we use to test a MongoDB instance controlled by AppleS with the smallest per-user I/O sequentiality, the  $\lambda_{\max}$  error is 4.46%.

**4.3.2  $\zeta$  Optimization.** We explore some classic feedback control approaches (e.g., Proportional-Integral (PI) and Proportional-Integral-Derivative (PID)), which are typically designed for high-precision control that leads to high complexity with less high-frequency robustness [63], especially when faced with highly-oscillating I/O-unfairness and latency-variability for the database under high-level user-parallelism. We currently adopt the low-cost additive-increase/multiplicative-decrease (AIMD) [8] feedback-control proven effective to bound I/O unfairness and latency variability (as evidenced by Fig. 9 and Fig. 11).

Specifically, AppleS performs fast timescale feedback control to maximize  $\zeta$  constrained by a user I/O unfairness upper bound  $U_f^o$  or a request latency variability threshold  $V_l^o$ . To be responsive to rapid changes of I/O unfairness and latency variations, AppleS monitors  $\zeta$  every scheduling round and

**Algorithm 2:**  $\zeta$  Optimization Algorithm

---

**Input:** A user I/O unfairness upper bound  $U_f^o$ , a request latency variability threshold  $V_l^o$

**Output:** The optimized  $\zeta$

```

1 /*Measure the latest user I/O unfairness.*/
2  $U_f \leftarrow \text{unfairness\_measure}(P_0)$ 
3 /*Measure the latest request latency variability.*/
4  $V_l \leftarrow \text{latencyvar\_measure}(P_1)$ 
5 /*If both thresholds are 0, keep the minimum  $\zeta$  for the
   lowest user I/O unfairness and latency variability.*/
6 if  $U_f^o = 0$  and  $V_l^o = 0$  then
7   return  $\zeta_{min}$ 
8 /*If  $U_f$  or  $V_l$  exceed their respective threshold,  $\zeta$  will
   exponentially decrease.*/
9 if  $(U_f^o > 0 \text{ and } U_f > U_f^o) \text{ or } (V_l^o > 0 \text{ and } V_l > V_l^o)$  then
10   if  $d * \zeta \geq \zeta_{min}$  then
11      $\zeta \leftarrow d * \zeta$ 
12   else
13      $\zeta \leftarrow \zeta_{min}$ 
14 /*If  $U_f$  or  $V_l$  are small enough to meet their thresholds,  $\zeta$ 
   will follow an additive increase.*/
15 if  $(U_f^o = 0 \text{ or } (U_f^o > 0 \text{ and } U_f < \eta_u * U_f^o)) \text{ and } (V_l^o = 0 \text{ or }
   (V_l^o > 0 \text{ and } V_l < \eta_v * V_l^o))$  then
16   if  $\zeta + u \leq \zeta_{max}$  then
17      $\zeta \leftarrow \zeta + u$ 
18   else
19      $\zeta \leftarrow \zeta_{max}$ 
20 return  $\zeta$ 

```

---

adjusts  $\zeta$  based on the basic principles of TCP congestion control [40], i.e., the AIMD feedback control algorithm [8]. Generally speaking, the fast-timescale feedback control of AppleS is also a kind of congestion control to maximize per-user request sequentiality afforded by the database under the matching point. Therefore, we borrow the basic idea of the TCP congestion control for the implementation of the  $\zeta$  optimization algorithm. Specifically, as shown in Algorithm 2, if the last increment of  $\zeta$  causes the measured user I/O unfairness or request latency variability to exceed their respective threshold,  $\zeta$  will be reduced by multiplying with  $d$  that must be less than 1 (set at 0.5, the same as the control parameter used by the TCP congestion control, by default). It will lead to exponential backoff of  $\zeta$  to rapidly respond to violations of upper bounds for user I/O unfairness and request latency variability (i.e.,  $U_f^o$  and  $V_l^o$ ). If both the measured user I/O unfairness and the latency variability are lower than their respective thresholds (i.e.,  $\eta_u * U_f^o$  and  $\eta_v * V_l^o$  where  $\eta_u$  and  $\eta_v$  are set at 0.9 by default), then it means that the current level of user I/O performance isolation allows sufficient room for optimizing  $\zeta$  for a stronger per-user request sequentiality. In

this case,  $\zeta$  is increased to exploit cache locality to optimize the I/O performance. If both  $U_f^o$  and  $V_l^o$  are set at 0, then it means that the lowest levels of user I/O unfairness and request latency variability are required and thus the minimum  $\zeta$  is returned.

## 5 Performance Evaluation

**Test Environment:** All the evaluation experiments are conducted on a dedicated rack of PowerEdge R630 servers. The storage server is equipped with a RAID-0 SSD array with five 800GB SATA MLC Solid State Drives, consolidating all the logical volumes for databases. The computing server is configured with 2 Intel Xeon E5-2650 processors, 64GB of RAM, a Broadcom NetXtreme II BCM57810 10Gb NIC and 4  $\times$  1TB SATA HDDs. All the servers are connected by a Dell N4032F switch with peak bandwidth of 10Gb.

**Workloads, Databases, I/O Schedulers and Resource Management:** We deploy a relational database, MySQL 8.0.15 [23], and a key-value store, MongoDB 3.6.0 [21], as representative databases to verify the effectiveness and robustness of AppleS on enforcing fair, stable and efficient I/O sharing for user-level QoS improvement. Specifically, we use the TPC-C benchmark [26] to establish multiple user connections to concurrently access a database consisting of 1,000 warehouses (for a total dataset size of 200GB) built on MySQL. Similarly, we run the YCSB Benchmark [32] on MongoDB with multiple user connections, each generating a Zipf distributed key-value request workload. User request workloads include different combinations of GET and SET, and are write-heavy (50% GET, 50% SET) unless otherwise noted, accessing the underlying MongoDB that stores a 150GB dataset. Since very small key-value objects (typically smaller than 1KB) are prevalent in enterprise-level stores [64], we set object size at 1KB for the MongoDB dataset. MySQL and MongoDB are each deployed on a computing server and their respective datasets are stored on their own logic volumes provided by a storage server. To verify the portability and compatibility, we evaluate AppleS under 9 I/O schedulers and across 4 versions of Linux kernels, as shown in Table 1. Since CFQ [14] may stall some requests in the queue needlessly and cause poor response times for MySQL [62] and BFQ [12] also exhibits poor performance for relational databases [6], we exclude these two schedulers from the MySQL evaluation. In addition, Split-AFQ, Split-Deadline, and Split-Token are officially customized for the Linux kernel 3.2.51 [66] and thus we adopt this kernel version for their assessments.

To compare with the widely adopted resource management mechanism, i.e., Cgroup(v2) [3], we install Cgroup(v2)-enabled CentOS 8.3 with the Linux kernel 5.10.10 to enable full functionality of Cgroup(v2), especially the support for thread-level CPU fairness enforcement. It's threaded controller can effectively regulate thread-level CPU contention

and guarantee thread-level CPU fair sharing. For CentOS 8.3, we install MySQL 8.0.23 and MongoDB 4.4.3 from their official repositories.

**Evaluation Metrics:** MySQL [23] and MongoDB [21], as different types of databases, focus on distinct workloads each of which can be characterized by disparate requirements for I/O performance quality of service and thus can have different evaluation metrics. As a representative relational database, MySQL mainly focuses on OLTP workloads (e.g., the TPC-C workload [26]) that demand high transaction throughput while MongoDB strives for highly responsive key-value processing that is more sensitive to tail latency. Nevertheless, both database types desire fair and efficient I/O sharing among concurrent users. According to their individual unique and common demands on I/O performance quality of service, we design the stated methodology for AppleS's evaluation on MySQL and MongoDB. Specifically, for the TPC-C workload, we use transactions per minute (TpmC) as the metric to assess its throughput while the transaction throughput in terms of *ops/s* is used for the YCSB [32]. In addition, we adopt user-level I/O unfairness (defined in Section 3.1) and the coefficient of variation (CV) [13] averaged across all users as metrics to evaluate user-experienced fairness and request latency variability, respectively.


**Baselines and Objectives:** Section 5.1 assesses the capability of AppleS to make a desirable tradeoff among transaction throughput, user I/O fairness and request-latency variability for the relational database under different levels of high user parallelism. Section 5.2 verifies the effectiveness of AppleS to enforce fair and stable I/O sharing across different numbers of concurrent users by making full use of I/O resources. Section 5.3 investigates AppleS' ability to reduce I/O provisioning for user QoS guarantees, e.g., enforcing the 99.9<sup>th</sup> percentile (i.e.,  $P_{99.9}$ ) tail latency SLOs.

OS	Linux kernel	Resource management or I/O scheduler
CentOS 7.3	3.2.51	Split-AFQ; Split-Deadline; Split-Token
CentOS 7.3	4.9.75	CFQ; Deadline; Noop
CentOS 7.3	5.0.5	mq-deadline; Kyber; BFQ
CentOS 8.3	5.10.10	Cgroup (v2) & mq-deadline


**Table 1.** The baseline systems and environments.

### 5.1 AppleS's Optimization for MySQL


In this evaluation, we assess how effectively AppleS optimizes MySQL in terms of the tradeoff among throughput, user I/O fairness and request-latency variability under high user parallelism. Based on the investigation in Section 2.1, a number of concurrent users higher than 64 can generate excessive OLTP workloads on the MySQL running on our testbed. We thus adopt 64 or more users in this assessment. Specifically, AppleS optimizes the scheduling-round I/O quota  $\zeta$  at the matching point to exploit potential spatial



**Figure 9.** The transaction throughput optimization for MySQL 8.0.15 under the constraints of user I/O unfairness and request-latency variability.




**Figure 10.** The breakdowns of CPU usage on the code-paths related to lock contention and cache conflict for MySQL 8.0.15.



**Figure 11.** Transaction throughput and unfairness under different schedulers for MySQL 8.0.15.


locality to further optimize the throughput performance, subject to two conditions, i.e., user I/O unfairness is less than or equal to 15% and the level of request-latency variability (by measuring the coefficient of variation (CV) [13]) is less than or equal to 90% of the CV of request latency measured for the baseline system (i.e., MySQL running on mq-deadline [22]). Based on the fast-timescale feedback control with 100ms-level monitoring for changes of user I/O unfairness and the CV of request latency, AppleS can effectively respond to each condition-violation event by quickly adjusting  $\zeta$ . As show in Fig. 9, user I/O unfairness and the CV of request latency can be effectively confined within the expected ranges except for the 512-user case in which the unfairness value is 17% and slightly higher than the target. The throughput improvement ranges from 9.8% to 28.1% when the number of users increases from 64 to 512. This is because AppleS effectively improves the efficiency of the buffer pool by increasing its hit ratio and cutting down the extra disk I/Os for futile page buffering. Also evidenced by Fig. 10, AppleS successfully hides the excessive user parallelism that, otherwise, would aggravate lock contentions and buffer cache conflicts, thus significantly reducing the CPU usage on the relevant code-paths.




**Figure 12.** AppleS’s effectiveness in reducing user I/O unfairness and request-latency variability, and increasing total throughput for MySQL 8.0.23 under Cgroup(v2)-enabled CentOS 8.3.

We further verify the effectiveness of AppleS in optimizing transaction throughput conditioned by the unfairness target of 10% under 7 I/O schedulers across 3 versions of Linux kernels respectively and compare the transaction throughput and unfairness measured with and without AppleS respectively, giving rise to the corresponding pair-wise comparisons. As demonstrated in Fig. 11, AppleS is able to improve the throughput performance for all the cases studied, with as much as 17.3% improvement for the mq-deadline case. Note that the unfairness values measured under AppleS across all the cases are below or very close to the required limit (i.e., 10%), improving over the cases without AppleS by  $4.6\times \sim 7.3\times$ . The highest unfairness incurred by AppleS is merely 11.7% for the AppleS+kyber case.


We also assess AppleS’s capability to cooperate with Cgroup(v2) [3] on MySQL 8.0.23. Although the existing Cgroup(v2) can enforce process/thread-level resource distribution hierarchically in a controlled and configurable manner for a subset of resource controllers (e.g., weight-based CPU controller), it is not yet capable of controlling thread-level fair I/O sharing [3]. Thus, MySQL user-level I/O behavior, which typically happens before the corresponding thread I/Os, is actually beyond the I/O control of Cgroup(v2). However, I/O processing conducted by MySQL’s work threads still rely on CPU cycles. Thus, we adopt Cgroup(v2)-enforced weight-based CPU controller to enforce thread-level CPU fair sharing among MySQL’s work threads. As shown in Fig. 12, Cgroup(v2)-enforced MySQL performs slightly better than the baseline (i.e., MySQL 8.0.23) on user-level I/O fairness and latency variability in some cases (e.g., the case of 512 user connections). However, as shown in Fig. 13, Cgroup(v2) cannot effectively alleviate the contentions on lock and buffer cache of MySQL, which are the dominant factors detrimental to I/O fairness, efficiency and latency performance under high user parallelism. In contrast, AppleS can effectively hide excessive user I/O parallelism from MySQL and its underlying I/O stack, and thus significantly alleviate lock contentions and cache conflicts within MySQL, as shown in Fig. 13. This explains why AppleS can significantly enhance user-level I/O fairness by up to  $31.8\times$  with a low latency variability, as well as improve throughput performance by up to 39.2%. More importantly, when combined with Cgroup(v2), AppleS is shown to be orthogonal and complementary to



**Figure 13.** The breakdowns of CPU usage on the code-paths related to lock contention and cache conflict for MySQL 8.0.23 under Cgroup(v2)-enabled CentOS 8.3.



**Figure 14.** AppleS’s effectiveness in reducing user I/O unfairness and request-latency variability, and increasing throughput for MongoDB 4.4.3 under Cgroup(v2)-enabled CentOS 8.3.




**Figure 15.** The breakdowns of CPU usage for MongoDB 4.4.3 under Cgroup(v2)-enabled CentOS 8.3.

Cgroup(v2);<sup>–</sup>’s CPU fair scheduling efforts on MySQL work threads, further improving AppleS’s advantages on enhancing user-level I/O fairness (up to  $34.5\times$ ) and throughput (up to 45.6%).

## 5.2 AppleS’s Optimization for MongoDB

In this section, we verify AppleS’s efficacy on MongoDB 4.4.3 by adopting the same Cgroup(v2) configuration as that for the MySQL 8.0.23 case (described in Section 5.1) so that thread-level CPU fair sharing is enabled for MongoDB. Compared with the OLTP workload under the relational database MySQL with complex transaction processing but slower I/O speed, YCSB workloads [32] accessing MongoDB can generate a higher key-value request rate that relies more on stable and adequate CPU cycle provisioning. However, thread-level contention on CPU resource can force some threads to forfeit their CPU shares and exhibit deceptive idleness [67], resulting in low utilization of CPU and thus slower I/O speed for the storage I/O stack. As shown in Fig. 14 and Fig. 15, Cgroup(v2)-enforced thread-level CPU fair sharing can help MongoDB’s work threads to expose their actual CPU demand to the OS scheduler and put a higher I/O pressure





**Figure 16.** Transaction throughput and unfairness under different schedulers for MongoDB 3.6.0.


on the storage (e.g., up to  $1.77\times$  I/O wait cpu share) than that under the baseline, leading to a higher throughput (up to 22.1% improvement). In addition, although Cgroup(v2) provides some benefit in reducing user-level I/O unfairness and latency variability in most cases, it fails to effectively regulate user-level I/Os. In contrast, AppleS can effectively hide excessive user parallelism from the I/O stack and significantly reduce the share of I/O wait CPU time by up to 71.3%. This indicates that more portions of the I/O journey will be under AppleS’s predictable control with user-level I/O fairness guarantee without sacrificing throughput, effectively preventing unpredictable and harmful I/O delay in the I/O stack. More importantly, AppleS can cooperate well with Cgroup (v2) to combine its efforts on CPU fair sharing and further improve throughput (by up to 28.6%) and user-level I/O fairness (by up to  $101.5\times$ ) with a low latency variability.

Further, we examine the effectiveness of AppleS for MongoDB [21] under different I/O schedulers. Specifically, we let MongoDB run on 9 I/O schedulers across 3 versions of Linux kernels, with and without AppleS, respectively. As shown in Figure 16, AppleS helps reduce unfairness by  $4.6\times \sim 39.6\times$  across all the 9 I/O schedulers. More importantly, the cost on throughput is negligibly small, with a maximum of 2.5% transaction throughput reduction in the Split-AFQ [66] case. AppleS helps improve throughput performance for some other I/O schedulers, e.g., CFQ [14], Deadline [16] and BFD [12]. Such improvements are likely to be attributed to AppleS’s optimization that effectively alleviates I/O blocking in the I/O stack by curbing queuing effects.


### 5.3 AppleS’s Optimization for User QoS

In this assessment, we verify the effectiveness of AppleS on reducing I/O provisioning for user QoS guarantee. AppleS is designed to export the interface for coordinating user-level fairness policy to other QoS guarantee systems, e.g., PSLO [49]. In this case, AppleS actually works as a user-level I/O scheduler for the database to provide fair, stable and efficient I/O scheduling that can be controlled by these QoS-guaranteed systems.

Specifically, we use PSLO as a QoS guarantee system for testing. PSLO can enforce any percentile tail latency SLOs



**Figure 17.** The time-trend curves of request-latency variability for the LS users under different schemes for different combinations of LS users and BE users.



**Figure 18.** Comparisons in the throughput allocated to the LS users and the throughput allocated to the BE users under different schemes for different combinations of LS users and BE users.

for consolidated virtual machines (VMs) by dynamically coordinating I/O resources allocated among VMs based on a measurement-based feedback control model [49]. We implement the PSLO control model for database users instead of VMs. In this way, PSLO can work with AppleS to guarantee the P99.9 tail latency SLO of 60ms for the latency-sensitive (LS) users accessing the database (e.g., MongoDB [21]) by coordinating user weights. Meanwhile, the majority of the users demand high throughput but run in a best-effort mode, especially in cloud computing systems. They are called best-effort (BE) users that often share the same database with the LS users. Thus, cloud service providers would like to use minimum I/O resources needed to meet the LS users’ tail latency targets so as to spare as much I/O resources as possible for the BE users. In our evaluation, we use three combinations of LS users and BE users to share a MongoDB. The number of the LS users is kept at 10 while that of the BE users are 90, 140 and 190. Each LS user has a P99.9 tail latency SLO of 60ms. As shown in Fig. 17, AppleS can effectively reduce request-latency variability for the LS users, which becomes more important when more BE users share the MongoDB. As evidenced in Fig. 18, AppleS can successfully help PSLO to reduce I/O provisioning by cutting down the throughput allocated to the LS users (up to 53.1% reduction) and improve the throughput of the BE users (up to  $7.6\times$  improvement).

## 5.4 Takeaways from the Evaluation

The first takeaway from the evaluation is that it is possible to tackle different performance bottlenecks by the same solution. As we observe from the above evaluation, MySQL [23] and MongoDB [21] experience distinctively different performance bottlenecks. For highly scalable OLTP transaction processing, which widely adopts concurrency control to guarantee data consistency and highly relies on buffer cache to avoid unnecessary disk accesses, MySQL's I/O bottlenecks largely reside in the database internals (e.g., buffer cache and lock mechanisms). These bottlenecks are traditionally addressed by improving database internal components. In contrast, MongoDB's major performance bottlenecks manifest themselves in the storage I/O stack beneath the database for parallel processing of key-value operations, which can be treated as a classic storage-specific performance problem. Nevertheless, AppleS directly tackles the major root causes for the I/O bottlenecks encountered by both MySQL and MongoDB, i.e., excessive concurrent user requests exceeding the I/O capacity of the database-I/O-stack "pipe", and thus is effective for both on improving their user I/O performance isolation with high I/O efficiency. The second takeaway is on the importance and challenges of how to design an optimization scheme to synergize with the existing scheduling/optimization mechanisms. The key is how to minimize the "mutual-cancelling" effect when multiple optimizations work together. AppleS resides in the upper stream of the database-I/O-stack pipe and treats all the underlying I/O optimizations as a black box. This design is based on the experimental observations that both the database and storage I/O stack cannot fairly and efficiently respond to highly scaled concurrent user requests due to excessive parallel request contention. AppleS regulates parallel I/Os and effectively alleviates concurrent I/O contention along the database-I/O-stack pipe, which enables the databases and their underlying I/O scheduling/optimization mechanisms to work more efficiently. Moreover, we also observe that AppleS can synergize with Cgroup to complementarily improve user-level I/O performance. This is possible largely because they schedule different types of resources, i.e., AppleS focuses on reshaping parallel user I/Os while Cgroup's threaded controller mainly handles thread-level CPU resource sharing. In other words, the more knowledge AppleS has of various optimizations within the database-I/O-stack pipe, the better will AppleS be able to synergize with the latter for further improvement, but at the likely cost of the former's portability/flexibility. It is a potential direction for our future research.

## 6 Conclusion

In this paper, we propose an application-level I/O scheduler (AppleS) to support user-level scheduling goals of I/O fairness, latency request variability and throughput for database

systems. AppleS can proactively hide excessive user parallelism according to the derived matching point and optimize user I/O quotas under SLO constraints. Our extensive evaluation, driven by MySQL and MongoDB, demonstrates that users of databases can greatly benefit from AppleS, resulting in significantly enhanced user-level I/O fairness by up to 40.6× with a low variability of request latency, as well as an improved throughput by up to 39.2%. In addition, AppleS can help other QoS guarantee systems reduce I/O provisioning and run in a more cost-effective fashion. For the future work, we will explore the potentials of the synergy between AppleS and various optimizations along the I/O stack by treating the underlying subsystems as white or gray boxes.

## 7 Acknowledgments

This work is supported in part by the US National Science Foundation grants CCF-1704504 and CNS-2008835. We thank the anonymous reviewers for their valuable comments. We are especially grateful to our shepherd, Neeraja Yadwadkar, for helping improve the paper's presentation.

## A Proof of Eq. 2

*Proof.* Consider a simplified execution model that approximates the actual AppleS execution process, i.e., parallel execution rounds of  $P$  requests each and sequential execution rounds of one request each. Let  $p$  and  $m$  denote the total numbers of requests in the parallel execution rounds and the sequential execution rounds, respectively. Then for  $N$  total requests scheduled,  $N = p + m$ . We further define  $S_k^{(q)}$  as the service time for the  $k^{th}$  request of the execution type  $q$ , for sequential execution  $q = 1$  or parallel execution  $q = P$ . Then the total service time for the requests in the parallel execution ( $q = P$ ) is  $T_p = \sum_{k=1}^p S_k^{(P)}$  and the total service time for the requests in the sequential execution ( $q = 1$ ) is  $T_s = \sum_{k=1}^m S_k^{(1)}$ . With this model, the parallel fraction,  $\alpha$ , can then be explicitly expressed as,  $\alpha = T_p/T$ , where  $T = T_p + T_s$ , the sum of the request execution times for all the requests.

Thus, the total execution time of all  $N$  requests can be expressed as,  $\alpha * T/P + (1 - \alpha) * T$ , and hence, we have  $\lambda(P) = N/[\alpha(P) * T/P + (1 - \alpha(P)) * T]$ . As  $P$  reduces to a certain value  $P_0$  ( $< C$ ),  $\alpha$  is expected to approach 1. Substituting  $P = P_0$  and  $\alpha(P_0) = 1$  into the equation, we have  $\lambda(P_0) = P_0 * N/T$ . Since  $\alpha(P_0) = 1$ , from Eq. 1, we have  $A = 1 - B * e^{-\gamma * P_0}$ , by substituting which into Eq. 1, and then Eqs. 1 and  $\lambda(P_0) = P_0 * N/T$  into  $\lambda(P) = N/[\alpha(P) * T/P + (1 - \alpha(P)) * T]$ , we finally have Eq. 2.

## References

- [1] 2014. Database/kernel community topic at collaboration summit 2014. <http://www.postgresql.org/message-id/20140310101537.GC10663@suse.de>.
- [2] 2015. Cgroups v2. <https://www.kernel.org/doc/Documentation/cgroup-v2.txt>.

- [3] 2015. Control Group v2. <https://www.kernel.org/doc/html/latest/admin-guide/cgroup-v2.html>.
- [4] 2016. Janus codebase. <https://github.com/NYU-NEWS/janus>.
- [5] 2019. The system call intercepting library. [https://github.com/pmem/syscall\\_intercept](https://github.com/pmem/syscall_intercept).
- [6] 2020. Linux 5.6 I/O Scheduler Benchmarks: None, Kyber, BFQ, MQ-Deadline. <https://www.phoronix.com/scan.php?page=article&item=linux-56-nvme&num=3>.
- [7] 2021. 300,000 IOPS per Volume with Ultra High Performance Oracle Cloud Infrastructure Block Volumes. <https://blogs.oracle.com/cloud-infrastructure/post/300000-iops-per-volume-with-ultra-high-performance-oracle-cloud-infrastructure-block-volumes>.
- [8] 2021. Additive increase/multiplicative decrease. [https://en.wikipedia.org/wiki/Additive\\_increase/multiplicative\\_decrease](https://en.wikipedia.org/wiki/Additive_increase/multiplicative_decrease).
- [9] 2021. fsync. <https://docs.mongodb.com/manual/reference/command/fsync/>.
- [10] 2021. Stability (learning theory). [https://en.wikipedia.org/wiki/Stability\\_\(learning\\_theory\)](https://en.wikipedia.org/wiki/Stability_(learning_theory)).
- [11] 2022. Approximation error. [https://en.wikipedia.org/wiki/Approximation\\_error](https://en.wikipedia.org/wiki/Approximation_error).
- [12] 2022. BFQ. <https://www.kernel.org/doc/html/latest/block/bfq-iosched.html>.
- [13] 2022. Coefficient of variation. [https://en.wikipedia.org/wiki/Coefficient\\_of\\_variation](https://en.wikipedia.org/wiki/Coefficient_of_variation).
- [14] 2022. Completely Fair Queuing. <https://www.kernel.org/doc/Documentation/block/cfq-iosched.txt>.
- [15] 2022. DB2. <http://www.ibm.com/software/db2/>.
- [16] 2022. Deadline. <https://www.kernel.org/doc/Documentation/block/deadline-iosched.txt>.
- [17] 2022. Google Earth. <https://www.google.com/earth/>.
- [18] 2022. Google Finance. <https://www.google.com/finance>.
- [19] 2022. InnoDB Locking. <https://dev.mysql.com/doc/refman/8.0/en/innodb-locking.html>.
- [20] 2022. Kyber. <https://www.kernel.org/doc/Documentation/block/kyber-iosched.txt>.
- [21] 2022. MongoDB. <http://www.mongodb.org/>.
- [22] 2022. mq-deadline. <https://github.com/torvalds/linux/blob/master/block/mq-deadline.c>.
- [23] 2022. MySQL. <http://www.mysql.com>.
- [24] 2022. Newton's method. [https://en.wikipedia.org/wiki/Newton's\\_method](https://en.wikipedia.org/wiki/Newton's_method).
- [25] 2022. Oracle. <https://www.oracle.com>.
- [26] 2022. TRANSACTION PROCESSING PERFORMANCE COUNCIL. The TPC-C home page. <http://www.tpc.org/tpcc/>.
- [27] 2022. Weighted round robin. [https://en.wikipedia.org/wiki/Weighted\\_round\\_robin](https://en.wikipedia.org/wiki/Weighted_round_robin).
- [28] Oana Balmau, Florin Dinu, Willy Zwaenepoel, Karan Gupta, Ravishankar Chandhiramoorthi, and Diego Didona. 2019. SILK: Preventing Latency Spikes in Log-Structured Merge Key-Value Stores. In *Proceedings of the USENIX Annual Technical Conference (ATC)*.
- [29] A. Bouch, A. Kuchinsky, and N. Bhatti. 2000. Quality is in the eye of the beholder: meeting users' requirements for Internet quality of service. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI'00) (The Hague, The Netherlands, April 2000)*.
- [30] Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wallach, Mike Burrows, Tushar Chandra, Andrew Fikes, and Robert E. Gruber. 2006. Bigtable: A Distributed Storage System for Structured Data. In *Proceedings of the Symposium on Operating Systems Design and Implementation (OSDI)*.
- [31] Xusheng Chen, Haoze Song, Jianyu Jiang, Chaoyi Ruan, Cheng Li, Sen Wang, Gong Zhang, Reynold Cheng, and Heming Cui. 2021. Achieving low tail-latency and high scalability for serializable transactions in edge computing. In *Proceedings of the 3th European conference on Computer systems (EuroSys)*.
- [32] B. F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears. 2010. Benchmarking cloud serving systems with YCSB. In *Proceedings of the ACM symposium on Cloud computing (SoCC)*.
- [33] S. S. Craciunas, C. M. Kirsch, and H. Röck. 2008. I/O Resource Management Through System Call Scheduling. *SIGOPS Oper. Syst. Rev.* 42, 5 (2008), 44–54.
- [34] Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin, Swami Sivasubramanian, Peter Vosshall, and Werner Vogels. 2007. Dynamo: Amazon's Highly Available Key-Value Store. In *Proceedings of the Symposium on Operating Systems Design and Implementation (OSDI)*.
- [35] Akon Dey, Alan Fekete, and Uwe Röhm. 2013. Scalable transactions across heterogeneous NoSQL key-value data stores. In *Proceedings of the VLDB Endowment*.
- [36] Dave Dice, Virendra J. Marathe, and Nir Shavit. 2014. Persistent unfairness arising from cache residency imbalance. In *Proceedings of the 26th ACM symposium on Parallelism in algorithms and architectures*.
- [37] Michael Ferdman, Almutaz Adileh, Onur Kocberber, Stavros Volos, Mohammad Alisafae, Djordje Jevdjic, Cansu Kaynak, Adrian Daniel Popescu, Anastasia Ailamaki, and Babak Falsai. 2012. Clearing the clouds: A study of emerging scale-out workloads on modern hardware. In *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*.
- [38] M. P. Grosvenor, M. Schwarzkopf, I. Gog, R. N. Watson, A. W. Moore, S. Hand, and J. Crowcroft. 2015. Queues don't matter when you can jump them!. In *Proceedings of the USENIX conference on Networked systems design and implementation (NSDI)*.
- [39] Takashi Horikawa. 2013. Latch-free data structures for DBMS: design, implementation, and evaluation. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data, SIGMOD'13*.
- [40] Van Jacobson. 1988. Congestion avoidance and control. In *SIGCOMM*.
- [41] K. Jang, J. Sherry, H. Ballani, and T. Moncaster. 2015. Silo: Predictable message latency in the cloud. In *SIGCOMM*.
- [42] Song Jiang and Xiaodong Zhang. 2002. LIRS: an efficient low interference recency set replacement policy to improve buffer cache performance. In *Proceedings of the International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*.
- [43] W. Jin, J. S. Chase, and J. Kaur. 2004. Interposed proportional sharing for a storage service utility. *ACM SIGMETRICS Performance Evaluation Review* 32, 1 (2004), 37–48.
- [44] Ryan Johnson, Ippokratis Pandis, and Anastasia Ailamaki. 2009. Improving OLTP scalability using speculative lock inheritance. In *Proceedings of the VLDB Endowment*.
- [45] T. Johnson and D. Shasha. 1994. 2Q: A low overhead high performance buffer management replacement algorithm. In *Proceedings of the VLDB Endowment*.
- [46] Hyungsoo Jung, Hyuck Han, Alan Fekete, Gernot Heiser, and Heon Y. Yeom. 2014. A scalable lock manager for multicores. *ACM Transactions on Database Systems (TODS)* 39, 4 (2014), 1–29.
- [47] M. Katevenis, S. Sidiropoulos, and C. Courcoubetis. 1991. Weighted round-robin cell multiplexing in a general-purpose ATM switch chip. *IEEE Journal on Selected Areas in Communications* 9, 8 (1991), 1265–1279.
- [48] H. Kim and S. Ahn. 2008. BPLRU: A buffer management scheme for improving random writes in flash storage. In *Proceedings of the conference on File and storage technologies (FAST)*.
- [49] N. Li, H. Jiang, D. Feng, and Z. Shi. 2016. PSLO: Enforcing the  $X^{th}$  Percentile Latency and Throughput SLOs for Consolidated VM Storage. In *Proceedings of the 3th European conference on Computer systems (EuroSys)*.
- [50] H. C. Lim, S. Babu, and J. S. Chase. 2010. Automated Control for Elastic Storage. In *IEEE International Conference on Autonomic Computing (ICAC)*.
- [51] Yanfei Lv, Bin Cui, Bingsheng He, and Xuexuan Chen. 2011. Operation-aware buffer management in flash-based systems. In *Proceedings of the*

- 2011 ACM SIGMOD International Conference on Management of Data, SIGMOD'11.
- [52] Nimrod Megiddo and Dharmendra S. Modha. 2003. ARC: A Self-Tuning, Low Overhead Replacement Cache. In *Proceedings of the conference on File and storage technologies (FAST)*.
- [53] Barzan Mozafari, Carlo Curino, Alekh Jindal, and Samuel Madden. 2013. Performance and resource modeling in highly-concurrent OLTP workloads. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data, SIGMOD'13*.
- [54] Barzan Mozafari, Carlo Curino, and Samuel Madden. 2013. DBSeer: Resource and Performance Prediction for Building a Next Generation Database Cloud. In *CIDR*.
- [55] L. N. Bairavasundaram, G. Soundararajan, V. Mathur, K. Voruganti, and S. Kleiman. 2011. Italian for beginners: the next steps for SLO-based management. In *Proceedings of the USENIX symposium on Hot Topics in Storage and File Systems (HotStorage)*.
- [56] Elizabeth J. O'Neil, Patrick E. O'Neil, and Gerhard Weikum. 1993. The LRU-K page replacement algorithm for database disk buffering. In *Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data, SIGMOD'93*.
- [57] Ippokratis Pandis, Ryan Johnson, Nikos Hardavellas, and Anastasia Ailamaki. 2010. Data-oriented transaction execution. In *Proceedings of the VLDB Endowment*.
- [58] Ippokratis Pandis, Pinar Tözün, Ryan Johnson, and Anastasia Ailamaki. 2011. PLP: page latch-free shared-everything OLTP. In *Proceedings of the VLDB Endowment*.
- [59] W. Reda, M. Canini, L. Suresh, D. Kostić, and S. Braithwaite. 2017. Rein: Taming Tail Latency in Key-Value Stores via Multiget Scheduling. In *Proceedings of the 3th European conference on Computer systems (EuroSys)*.
- [60] Kun Ren, Jose M. Faleiro, and Daniel J. Abadi. 2016. Design principles for scaling multi-core oltp under high contention. In *Proceedings of the 2016 ACM SIGMOD International Conference on Management of Data, SIGMOD'16*.
- [61] Kun Ren, Alexander Thomson, and Daniel J. Abadi. 2012. Lightweight locking for main memory database systems. In *Proceedings of the VLDB Endowment*.
- [62] Baron Schwartz, Peter Zaitsev, and Vadim Tkachenko. 2012. *High performance MySQL: optimization, backups, and replication*. O'Reilly Media, Inc., Sebastopol, CA.
- [63] Ramon Vilanova, Víctor M. Alfaro, and Orlando Arrieta. 2012. Robustness in PID control. *PID control in the third millennium* (2012), 113–145.
- [64] Xingbo Wu, Yuehai Xu, Zili Shao, and Song Jiang. 2015. LSM-trie: An LSM-tree-based Ultra-Large Key-Value Store for Small Data. In *Proceedings of the USENIX Annual Technical Conference (ATC)*.
- [65] Cong Yan and Alvin Cheung. 2016. Leveraging lock contention to improve OLTP application performance. In *Proceedings of the VLDB Endowment*.
- [66] S. Yang, T. Harter, N. Agrawal, S. S. Kowsalya, A. Krishnamurthy, S. Al-Kiswani, R. T. Kaushik, A. C. Arpaci-Dusseau, and R. H. Arpaci-Dusseau. 2015. Split-Level I/O Scheduling. In *Proceedings of the ACM Symposium on Operating Systems Principles (SOSP)*.
- [67] Yong Zhao, Kun Suo, Xiaofeng Wu, Jia Rao, Song Wu, and Hai Jin. 2019. Preemptive multi-queue fair queuing. In *Proceedings of the ACM International Symposium on High Performance Distributed Computing (HPDC)*.
- [68] T. Zhu, D. S. Berger, and M. Harchol-Balter. 2016. SNC-Meister: Admitting More Tenants with Tail Latency SLOs. In *Proceedings of the ACM symposium on Cloud computing (SoCC)*.
- [69] T. Zhu, A. Tumanov, M. A. Kozuch, M. Harchol-Balter, and G. R. Ganger. 2014. PriorityMeister: Tail Latency QoS for Shared Networked Storage. In *Proceedings of the ACM symposium on Cloud computing (SoCC)*.



## B Artifact Appendix

### B.1 Abstract

*AppleS aims to improve database scalability by delivering the right amount and pattern of user parallel I/O requests to the database system under excessive user parallelism, aligning with the concurrency supported by the database and its underlying I/O stack. In doing so, AppleS improves user-level I/O performance in terms of user-level I/O fairness, throughput and latency stability. Implemented as a user-space module based on system call interception, AppleS is compatible with and portable to different types/versions of databases, different versions of OS kernels and their resource management tools, e.g., Cgroups.*

### B.2 Description & Requirements

**B.2.1 How to access.** The source code of the AppleS artifact is now available at [https://github.com/NingBellWind/AppleS\\_Artifact](https://github.com/NingBellWind/AppleS_Artifact)<sup>4</sup> and all the necessary instructions for the experimental environment setup can also be found there.

**B.2.2 Hardware dependencies.** There are no hardware dependencies. B.3.1 provides guidance to constructing the physical testbed in such a way that the expected excessive user parallelism can be effectively reproduced under the given benchmarks.

**B.2.3 Software dependencies.** AppleS is implemented based on a `syscall_intercept` library.

**B.2.4 Benchmarks.** Please refer to the 2<sup>nd</sup> paragraph of Section 5 of the AppleS paper.

### B.3 Set-up

**B.3.1 Hardware setup.** The physical testbed consists of two PowerEdge R630 servers, a computing server and a storage server. The former is used to run database instances, the AppleS artifact, and benchmarks while the database files reside on the latter. Specifically, the computing server is configured with 2 Intel Xeon E5-2650 V4 processors, 64GB of RAM, a Broadcom NetXtreme II BCM57810 10Gb NIC, and 2 × 1TB SATA HDDs while the storage server is equipped with 2 Intel Xeon E5-2603 V4 processors, 64GB of RAM, a Broadcom NetXtreme II BCM57810 10Gb NIC, and a RAID-0 SSD array with five 800GB SATA MLC SSDs, consolidating all the logical volumes (LVs) (formatted as Ext4) for databases. All the servers are connected by a Dell N4032F switch with a peak bandwidth of 10Gb.

**B.3.2 OSes setup.** For the computing server, please use CentOS 8.3 with the Linux kernel 5.10.10 that should be enabled for full functionality of Cgroup(v2) while CentOS 7.3 with default Linux kernel can be installed on the storage server. All the OSs should enable the iSCSI protocol and install the necessary

development tools (i.e., check "File and Storage Server" and "Development Tools" when installing the OS).

**B.3.3 Database setup.** Two types of databases, i.e., MySQL and MongoDB, are used in this setup. They include two versions used in the AppleS evaluation, i.e., MySQL 8.0.23 and MongoDB 4.4.3. All the installation instructions for databases and their configurations can be found at [https://github.com/NingBellWind/AppleS\\_Artifact](https://github.com/NingBellWind/AppleS_Artifact).

**B.3.4 Storage setup.** Two LVs should be created on the storage server, one is for the MySQL database storage while the other is used for the MongoDB database. Based on the estimated size of data sets for benchmarks, the size of each LV should be 512GB or larger.

**B.3.5 Software dependencies.** Before compiling AppleS, `syscall_intercept` library should be first installed. Its source code and installation instructions are available at [https://github.com/pmem/syscall\\_intercept](https://github.com/pmem/syscall_intercept).

**B.3.6 AppleS setup.** AppleS setup consists two steps, i.e., compilation and configuration. First, the source code and the compiled `syscall_intercept` library are deployed in a source directory. And then the AppleS artifact can be compiled. Second, configure AppleS for the target database by the file `apples_configuration.txt`, which is deployed in the same directory of the database executable file (i.e., running directory). After that, one can manually start the database instance running with AppleS loaded by `LD_PRELOAD`.

### B.4 Evaluation workflow

This section includes two experiments that are conducted to evaluate AppleS's effectiveness on the two databases, MySQL and MongoDB, under excessive user parallelism and to validate the AppleS paper's key results and claims.

**B.4.1 Major Claims.** The major claims and key results made in the AppleS paper are listed as follows, noting that references in square brackets refer to the AppleS paper:

- (C1): AppleS can significantly enhance user-level I/O fairness by up to 31.8× with a low latency variability, as well as improve throughput performance by up to 39.2% for MySQL 8.0.23. This is proven by the experiment (E1) described in [Section 5.1] whose results are illustrated in [Fig. 12].
- (C2): AppleS can cooperate well with Cgroup (v2) to combine its effects on CPU fair sharing and further improve throughput (by up to 28.6%) and user-level I/O fairness (by up to 101.5×) with a low latency variability for MongoDB 4.4.3. This is proven by the experiments (E2) described in [Section 5.2] whose results are illustrated in [Fig. 14].
- (C3): AppleS is portable to different types/versions of databases (i.e., MongoDB and MySQL). This is proven by the experiments (E1 and E2) described in [Section 3.4].

<sup>4</sup>DOI: 10.5281/zenodo.6336004

- (C4): AppleS is orthogonal to existing kernel-level resource management tools (e.g., Cgroups). This is proven by the experiment (E2) described in [Section 3.4].

**B.4.2 Experiments.** This section introduces the following two experiments (i.e., E1 and E2):

*Experiment (E1): [MySQL 8.0.23+CentOS 8.3+Kernel 5.10.10] [30 human-minutes + 1 compute-hour]: E1 aims to assess AppleS's capability to improve throughput for MySQL 8.0.23 with significantly enhanced user-level fairness and a low latency variability.*

*[How to] E1 consists of 3 steps, i.e., baseline measurement, P optimization, and measurement under the AppleS control. Step 1: conduct 3 runs of TPC-C workload with 256 concurrent connections accessing MySQL 8.0.23 running with disabled AppleS, which only tracks user I/O statistics and records them in log files. Step 2: enable AppleS to conduct P optimization by setting "opt" at 1 and "disable" at 0. The P optimization is only required to run once for a specific database system and lasts about 6 minutes. Step 3: almost the same with Step 1 except for enabling AppleS by setting "disable" at 0. And then, one can compare between the measures obtained under the baseline and the AppleS-controlled case. The expected outcome would be over 20% throughput improvement, over 10× user I/O fairness enhancement, and over 5× lower latency variability.*

*[Preparation] Create the database and load data for the TPC-C benchmark, which is only required to do once and lasts about 18 hours. Set the disk scheduler for the block device where the database files reside on as "mq-deadline".*

*[Execution] Configure apples\_configuration.txt for each step and deploy it in running directory before starting the step.*

*[Results] Each run lasts 360 seconds, including 60-seconds ramp-up time and 300-seconds running time. The throughput measured under different cases can be directly observed from the benchmark results while CV of request latency and user I/O unfairness can be collected from the last two columns*

*of the  $a^{th}$  line ( $a = \text{running\_time}/\text{control\_interval}$ , e.g., if  $\text{running\_time}=300$  seconds and  $\text{control\_interval}=0.1$  second,  $a = 3000$ ) in `bslo_cmd_log1.txt`. Each measure is the average over the three runs.*

*Experiment (E2): [MongoDB 4.4.3+Cgroup (V2)] [30 human-minutes + 1 compute-hour]: E2 aims to verify AppleS's capability to cooperate with Cgroup (V2) for MongoDB 4.4.3 to gain further improvement on user-level fairness and latency variability.*

*[How to] E2 consists of 5 steps, i.e., baseline measurement, P optimization, measure under the AppleS control, measure under the control of Cgroup (V2), and measure under the control of both AppleS and Cgroup. Step 1: conduct 3 runs of YCSB workload with 512 concurrent connections accessing MongoDB 4.4.3 running with disabled AppleS. Step 2: the same with the Step 2 of E1. Step 3: almost the same with Step 1 except for enabling AppleS by setting "disable" at 0. Step 4: almost the same with Step 1 except for setting /cgroup2/cg1/cgroup.procs as the PID of MongoDB to enable the Cgroup (V2) control. Step 5: almost the same with Step 4 except for enabling AppleS. And then, one can compare among the measures obtained under the baseline, the AppleS-controlled case, the Cgroup-controlled case, and the dual-controlled case. The expected outcome is that AppleS can work with Cgroup to achieve a high user-level fairness (over 30× improvement than the baseline) and low latency variability (over 2× improvement than the baseline).*

*[Preparation] Create the database and load data for the YCSB benchmark, which is only required to do once and lasts about 16 hours. Set the disk scheduler for the block device where the database files reside on as "mq-deadline".*

*[Execution] Please refer to E1. In addition, configure Cgroup settings if one needs its control.*

*[Results] Please refer to E1.*

*One can find more instructions and details at [https://github.com/NingBellWind/AppleS\\_Artifact](https://github.com/NingBellWind/AppleS_Artifact).*