

# Async-fork: Mitigating Query Latency Spikes Incurred by the Fork-based Snapshot Mechanism from the OS Level

Pu Pang  
Shanghai Jiao Tong University  
Alibaba Group  
avengerispp@sjtu.edu.cn

Quan Chen  
Shanghai Jiao Tong University  
chen-quan@sjtu.edu.cn

Yu Xu  
Alibaba Group  
qiyu.xy@alibaba-inc.com

Xiyu Wang  
Alibaba Group  
xiyu.wxy@alibaba-inc.com

Yong Yang  
Alibaba Group  
zhiche.yy@alibaba-inc.com

Gang Deng  
Alibaba Group  
denggang.dg@alibaba-inc.com

Shixuan Sun  
National University of Singapore  
sunsx@comp.nus.edu.sg

Hongbo Yao  
Alibaba Group  
yuanzhi.yhb@alibaba-inc.com

Zheng Liu  
Alibaba Group  
wenqing.lz@alibaba-inc.com

Tao Ma  
Alibaba Group  
boyu.mt@alibaba-inc.com

Kaihao Bai  
Shanghai Jiao Tong University  
Alibaba Group  
asbaikaihao@sjtu.edu.cn

Bo Liu  
Shanghai Jiao Tong University  
boliu98@sjtu.edu.cn

Zhengheng Wang  
Alibaba Group  
zhengheng.wzh@alibaba-inc.com

Zhuo Song  
Alibaba Group, SJTU  
songzhuo.sz@alibaba-inc.com

Minyi Guo  
Shanghai Jiao Tong University  
myguo@sjtu.edu.cn

## ABSTRACT

In-memory key-value stores (IMKVSes) serve many online applications. They generally adopt the fork-based snapshot mechanism to support data backup. However, this method can result in query latency spikes because the engine is out-of-service for queries during the snapshot. In contrast to existing research optimizing snapshot algorithms, we address the problem from the operating system (OS) level, while keeping the data persistent mechanism in IMKVSes unchanged. Specifically, we first study the impact of the fork operation on query latency. Based on findings in the study, we propose Async-fork, which performs the fork operation asynchronously to reduce the out-of-service time of the engine. Async-fork is implemented in the Linux kernel and deployed into the online Redis database in public clouds. Our experiment results show that Async-fork can significantly reduce the tail latency of queries during the snapshot.

### PVLDB Reference Format:

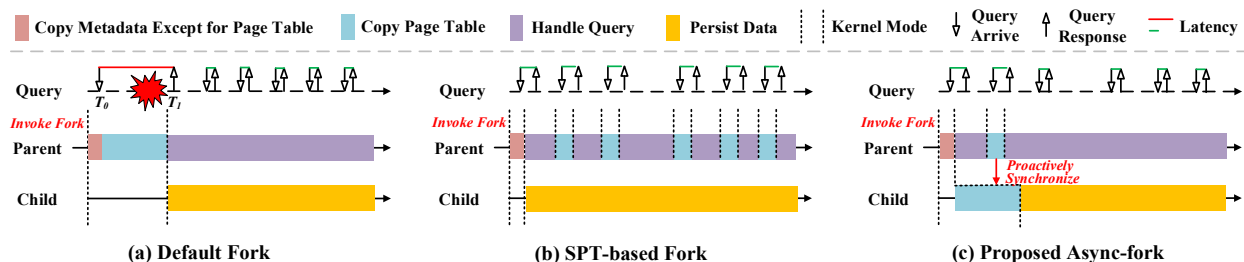
Pu Pang, Gang Deng, Kaihao Bai, Quan Chen, Shixuan Sun, Bo Liu, Yu Xu, Hongbo Yao, Zhengheng Wang, Xiyu Wang, Zheng Liu, Zhuo Song, Yong Yang, Tao Ma, and Minyi Guo. Async-fork: Mitigating Query Latency Spikes Incurred by the Fork-based Snapshot Mechanism from the OS Level. PVLDB, 16(5): 1033 - 1045, 2023.  
doi:10.14778/3579075.3579079

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing [info@vldb.org](mailto:info@vldb.org). Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.  
Proceedings of the VLDB Endowment, Vol. 16, No. 5 ISSN 2150-8097.  
doi:10.14778/3579075.3579079

## 1 INTRODUCTION

In-memory key-value stores (IMKVSes) are widely used in real-world applications, especially online services (e.g., e-commerce and social network), because of their ultra-fast query processing speed. For example, Memcached [49], Redis [42], KeyDB [24] and their variants [6, 11, 18, 23] have been deployed in production environments of big internet companies such as Facebook, Amazon and Twitter. As all data resides in memory, the data persistent function is a key feature of IMKVSes for data backup.

A common data persistent approach is to take a point-in-time snapshot of the in-memory data with the system call *fork* and dump the snapshot into the file system. Figure 1(a) gives an example of the fork-based snapshot method. In the beginning, the storage engine (the parent process) invokes *fork* to create a child process. As *fork* creates a new process by duplicating the parent process, the child process will hold the same data as the parent. Thus, we can ask the child process to write the data into a file in the background but keep the parent process continuous processing queries. Although the storage engine delegates the heavy IO task to the child process, the fork-based snapshot method can incur latency spikes [31, 52]. Specifically, queries arriving during the period of taking a snapshot (from the start of *fork* to the end of persisting data) can have a long latency because the storage engine runs into the kernel mode and is out-of-service for queries. For example, the query arrives in time  $T_0$  in Figure 1(a). For brevity, queries arriving during the period of taking a snapshot are called *snapshot queries*, while the others are called *normal queries*. In general, IMKVSes take a snapshot periodically (e.g., Redis by default takes a snapshot every 60 seconds if at least 10000 records are modified [47, 59]). Consequently, latency spikes for snapshot queries are not rare.



**Figure 1: The workflow of the parent and child process with (a) default fork, (b) shared page table (SPT)-based fork, (c) the proposed Async-fork in the snapshot procedure.**

To solve the problem, researchers proposed a variety of snapshot algorithms such as Copy-on-Update [9, 32], Zigzag [10] and Ping-Pong [10]. These algorithms focus on lowering the cost of taking a snapshot by optimizing the occasion of copying data and reducing the amount of data copied. However, a recent study [31] finds that the performance of the fork-based method is generally competitive with these advanced methods, and even outperforms them for write-intensive workloads. Additionally, the implementation of the fork-based method is very simple and requires a small engineering effort. As such, industrial IMKVSes do not adopt these snapshot algorithms from academia and keep the fork-based as the built-in data persistent approach. As a result, IMKVSes still have latency spikes caused by the fork operation, which can harm the service quality of online applications.

In contrast to existing research focusing on optimizing snapshot algorithms, we propose a straightforward approach to address the problem. Specifically, we optimize the fork operation to reduce the long latency incurred by the fork-based snapshot from the operating system (OS) level, while keeping the data persistent mechanism in IMKVSes unchanged.

We first conduct an in-depth study to reveal the impact of the fork operation on query latency. Our profiling results on the default fork operation show that the overhead of *fork* results in a long tail latency (up to hundreds of milliseconds) for snapshot queries, and copying the page table dominates the execution time of *fork*. This motivates us to further investigate the impact of advanced techniques [13, 14, 17, 34, 59] that can accelerate the fork operation. We consider two optimization strategies, the huge page [13, 14] and the shared page table [17, 34, 59]. As the huge page can degrade the performance of IMKVSes [27, 44], our profiling only involves the shared page table-based *fork* (SPT-based *fork*) [59], which is the latest method. The SPT-based *fork* proposed to share the page table between the parent process and the child process in a copy-on-write (CoW) manner to reduce the cost of the fork operation. However, the CoW can frequently interrupt the parent process as shown in Figure 1(b). Our profiling results show that frequent interruption can incur non-negligible overhead for snapshot queries although the SPT-based *fork* significantly reduces the tail latency compared with the default fork operation. Moreover, our analysis finds that the shared page table introduces the data leakage vulnerability, which potentially leads to an inconsistent snapshot. Thus, both the huge table and shared page table techniques cannot be applied in this scenario.

Motivated by these findings, we propose Async-fork to mitigate the latency spikes for snapshot queries by optimizing the fork operation. Figure 1(c) demonstrates the general idea. As copying the page table dominates the cost of the fork operation, Async-fork offloads this workload from the parent process to the child process to reduce the duration that the parent process runs into the kernel mode. This design also ensures that both the parent and child processes have an exclusive page table to avoid the data leakage vulnerability caused by the shared page table.

However, it is far from trivial to achieve in design since the asynchronization operations of the two processes on the page table can result in an inconsistent snapshot, i.e., the parent process may modify the page table, while the copy operation of the child process is in process. To address the problem, we design the proactive synchronization technique. This technique enables the parent process to detect all modifications (including that triggered by either users or OS) to the page table. If the parent process detects that some page table entries will be modified and these entries are not copied, then it will proactively copy them to the child process. Otherwise, these entries must have been copied to the child process and the parent process will directly modify them. In this way, the proactive synchronization technique keeps the snapshot consistent and reduces the number of interruptions to the parent process compared with SPT-based *fork*. Additionally, we parallelize the copy operation of the child process to further accelerate Async-fork.

We implement Async-fork in the Linux kernel (both x86 and ARM64). Async-fork is integrated into the OS and transparent to IMKVSes. The technique is also deployed in the online Redis databases in public Clouds<sup>1</sup>. Despite that, we conduct experiments on our local machine for the purpose of test flexibility. In the experiments, we select two popular IMKVSes, Redis and KeyDB, and use the Redis benchmark [41] and Memtier benchmark [46]. The database instance size is varied from 1GB, 2GB, and 4GB ... to 64GB. Although the SPT-based *fork* [59] may lead to data leakage, our experiment involves this method for comparison purposes because of its efficiency. For Redis, our experiment results show that 1) compared with the default *fork*, Async-fork reduces the 99%-ile latency of snapshot queries by 17.57% (from 0.074ms to 0.061ms) on 1GB instance, 81.76% (from 0.435ms to 0.079ms) on 8GB instance and 99.84% (from 991.9ms to 1.5ms) on 64GB instance; and 2) compared

<sup>1</sup><https://www.alibabacloud.com/product/apsaradb-for-redis>. Last accessed on 2022/11/13.




Figure 2: The organization of the page table.

with the latest SPT-based *fork*, Async-fork reduces the 99%-ile latency of snapshot queries by 2.87% on 1GB instance, 39.73% on 8GB instance and 61.97% on 64GB instance. We obtain similar results on KeyDB. These results demonstrate the efficacy of the technique proposed in this paper, especially for the large instances that can lead to long latency.

In summary, we make the following contributions in this paper.

- We conduct an in-depth study of the impact of the fork operations on the latency of snapshot queries in IMKVSes.
- We propose Async-fork that can mitigate the long latency of snapshot queries from the OS level, which is orthogonal to existing research on the problem.
- The technique is implemented in the Linux kernel (both x86 and ARM64) and deployed in the online Redis database in public clouds.
- We conduct extensive experiments with Redis and KeyDB to evaluate the efficacy of the proposed techniques.

## 2 BACKGROUND

In this section, we first introduce the preliminaries and then discuss the related work.

### 2.1 Preliminary

We first briefly review two operating system concepts, *virtual memory* and *fork* that are closely related to this work. As our technique is implemented and deployed in Linux, we introduce these concepts in the context of Linux. Then, we discuss the use cases of *fork* in databases.

**Virtual Memory.** As an effective approach to managing hardware memory resources, virtual memory is widely used in modern operating systems. A process has its own virtual memory space, which is organized into a set of virtual memory areas (VMA). Each VMA describes a continuous area in the virtual memory space. The page table is the data structure used to map the virtual memory space to the physical memory. It consists of a collection of page table entries (PTE), each of which maintains the virtual-to-physical address translation information and access permissions. A VMA corresponds to multiple PTEs.

Figure 2 shows an example of the page table. To reduce the memory cost, the page table is stored as a multi-level radix tree in which PTEs locate in leaf nodes (i.e., PTE Level in Figure 2).

The part in the area marked with the dashed line corresponds to VMA<sub>0</sub>. The tree at most has five levels. From top to bottom, they are the page global directory (PGD) level, the P4D level, the page upper directory (PUD) level, the page middle directory (PMD) level and the PTE level. As P4D is generally disabled, we focus on the other four levels in this paper. Except for the PTE level, an entry stores the physical address of a page while this page is used as the next-level node (table). With the page size setting to 4KB, a table in each level contains 512 entries. Given a VMA, “VMA’s PTEs” refers to the PTEs corresponding to the VMA and “VMA’s PMDs” is the set of PMD entries that are parents of these PTEs in the tree. For example, PMD<sub>0</sub> and PMD<sub>1</sub> belong to VMA<sub>0</sub>’s PMDs in Figure 2.

**Fork Operation.** *Fork* is a system call that creates a new process by duplicating the calling process [1]. Both processes have separate memory spaces. The new (resp. calling) process is called the child (resp. parent) process. To accelerate the operation, Linux implements *fork* with the copy-on-write (CoW) strategy. Specifically, while invoking *fork*, the parent process runs in the kernel mode and copies the metadata (e.g. VMAs, the page table, file descriptors, and signals) to the child process. The PTEs of both parent and child are set to write-protected. After that, the process that first modifies a write-protected page triggers a page fault, which leads to the copy of the page. In a word, benefiting from CoW, *fork* copies the metadata only.

**Use Cases of *fork* in Database.** *Fork* has a number of database use cases because it can easily and efficiently create a snapshot of in-memory data, the consistency of which is guaranteed by the OS. In general, these cases can be categorized into two classes based on the usage of the snapshot. First, use *fork* to delegate dedicated tasks, which have expensive IO or computation costs, to a child process without blocking the service of the parent process. MDC [40] uses *fork* to record checkpoints for in-memory databases. Redis uses *fork* to conduct *log rewriting* [45] that optimizes the Append Only File (AOF). FlurryDB [35] proposes to create replica based on *fork* in distributed environments. Second, use *fork* to create snapshots to support concurrent transaction processing because *fork* provides the snapshot isolation between processes. HyPer [29] proposes to evaluate hybrid OLTP and OLAP queries based on snapshots created by *fork*. AnKer [50] designs a fine-grained snapshot mechanism to support MVCC. In particular, AnKer takes a partial snapshot of in-memory data by co-designing the database engine and the system call *fork*. Different from our research on accelerating the fork operation from the OS level, AnKer focuses on optimizing which in-memory data should be captured by the snapshot.

All these use cases can potentially benefit from Async-fork because 1) they can encounter the query latency spike problem incurred by the fork operation; and 2) Async-fork can accelerate the snapshot creation. This paper focuses on the scenario that uses *fork* to take a point-in-time snapshot of in-memory data to persist the data [26, 45]. In particular, the storage engine calls *fork* to create a child process that holds the same data as it. Then, the child process writes the data to the hard drive, while the storage engine can continue to serve users’ queries. Although the storage engine delegates the data dump task to the child process, it will be out-of-service for queries during the invocation of *fork* because it runs into the kernel mode. We are particularly interested in this scenario because the IMKVS is one of the most important services in public

cloud and popular IMKVSes (e.g., Redis and KeyDB) [51] use this mechanism to persist data. Consequently, these stores encounter serious a query latency spike problem (see Section 3), while they are generally used in mission-critical applications that have a rigid latency constraint. We also evaluate the effectiveness of Async-fork on log rewriting in Redis [45]. The experiment results are presented in the technical report [37].

*Remarks.* Instead of developing a general-purpose solution to replace the default *fork* in the OS, the goal of Async-fork is to provide an efficient fork operation for the scenarios (especially for IMKVSes) where 1) the applications are memory-intensive, and 2) the parent process is latency-sensitive. Our design allows Async-fork and the default *fork* to run in parallel in the OS. Users can easily choose the fork method used in applications (see Section 5.2).

## 2.2 Related Work

Consistent snapshot is essential for in-memory databases to support backup and disaster recovery [4, 16, 31, 56]. Some consistent snapshot mechanisms have been proposed to trade off throughput, latency, and memory footprint [8–10, 31, 32]. Naive snapshot [8] blocked the storage engine until a deep copy of all the in-memory data is created, which is not suitable for IMKVSes in which the latency is critical. There are also some non-blocking snapshot mechanisms. Copy-on-Update [9, 32] proposed to create a shadow copy of the in-memory data; the storage engine is free to access any data but create a deep copy when updates it for the first time. Note that, the fork-based snapshot is a Copy-on-Update variant that leverages the operating system. Some other mechanisms used multi-version concurrency control (MVCC) [7] to keep multiple versions of in-memory data. Zigzag [10] maintained another untouched copy of the in-memory data and introduced metadata bits to indicate which copy the store engine should read from or write to. Based on Zigzag, Ping-Pong [10] maintained three versions of the data to lower the cost of managing metadata bits. Hourglass and Piggyback [31] were developed by combining Zigzag and Ping-Pong.

Although fork-based snapshot results in long latency during snapshot process, popular industrial IMKVSes (Redis and KeyDB) still adopts fork-based snapshot for two reasons: 1) *fork* provides a simple engineering implementation for consistent snapshot, while it requires great efforts to integrate the above approaches into the IMKVS. 2) None of the above approaches completely outperform the fork-based snapshot in write-intensive workloads [31]. For example, Ping-Pong and Hourglass can mitigate the latency spikes, while Ping-Pong incurs 3x memory footprint and Hourglass results in higher latency during normal operation. This work resolves the latency spikes of the fork-based snapshot, while keeping its original superiorities.

Previous work [32, 40] noted that the memory footprint increases during snapshot process due to the CoW strategy. MDC [40] proposed to release the pages that have been persisted as soon as possible. AnKer [50] introduced a fine-grained version of *fork* to take partial snapshot when databases do not need to persist all the data. CCoW [22] optimized the CoW mechanism based on the spatial locality of memory access. It prioritizes the copy for high-locality memory regions to improve the performance on write-intensive workloads. Async-fork is orthogonal and complementary to them.




Figure 3: The execution time of *fork* and the percentage of the time of copying the page table in *fork*.

There is also other approach to persist data in IMKVSes. Redis and KeyDB use Append Only File (AOF) [26, 45] to log every write operation received by the storage engine, that will be played again after the database reboots to reconstruct the original dataset. The snapshot and AOF are complementary, and it is recommended to enable both of them simultaneously in IMKVSes.

## 3 MOTIVATION

In this section, we present our profiling results to demonstrate the impact of the fork operation on the query latency. We first evaluate the performance of *fork* for taking the snapshot to pinpoint the key performance factors. We then reveal the impact of the fork operation on query latency. In addition to the default *fork* in Linux, our profiling involves the state-of-the-art approach [59] of optimizing the fork operation. Lastly, we summarize our findings according to the profiling results.

**Profiling Setting.** We use Redis benchmark [41] to study the performance in the experiments. The detailed configuration of the test machine is introduced in Section 6.1. The experiment reports the latency of a query, that is the elapsed time between the time point that the client issues the query and that the client receives the response. In the experiment, we enhance the benchmark by generating queries in the open-loop mode [48, 58]. This enhancement sends commands to the server without waiting for replies to previous queries to simulate real-world environments. The database instance size is varied from 1GB to 64GB. By default, Redis takes one snapshot per 60 seconds if at least 10000 keys changed. In order to measure the impact of the fork operation accurately, we execute the BGSAVE command to trigger the operation of taking a snapshot. We classify queries into two groups, *normal* and *snapshot*, based on their arrival time. The snapshot queries are the queries arriving during the period of taking the snapshot (i.e., from the invocation of *fork* until the end of persisting the in-memory data), while the others are normal queries. We measure the 99%-ile (p99) latency and the maximum latency of normal and snapshot queries, respectively. The two latencies greatly impact user experiences and are often used to measure the performance of user-facing databases [28, 31]. We repeat each experiment five times and report the average value.

### 3.1 Performance of *fork* for Taking a Snapshot

Figure 3 presents the execution time of *fork* and the time of copying the page table in the fork operation. We can see that the execution time grows roughly linearly with the instance size increasing from 1GB to 64GB. The 1GB instance takes less than 10 ms, while the

64GB instance takes more than 600 ms. We can also see that the time of the page table copy dominates the execution time. Particularly, the copy operation takes over 97% percentage of the execution time on all test cases and up to 99.93% percentage on the 64 GB instance. Without loss of generality, we measure the detailed metrics of copying the page table on the 8 GB instance to further study the copy operation.

In the experiment, the page table of the 8GB instance has one PGD entry, eight PUDs,  $2^{12}$  PMDs and  $2^{21}$  PTEs. Overall, *fork* copies the table level-by-level from top to bottom along the radix tree. The copy of one PGD/PUD/PMD entry requires to apply for a page to store its children and initialize the page. The operation takes around 500 ns. Thus, the copy of the  $2^{12}$  PMDs takes around 2 ms, while the overhead of copying PGDs and PUDs is trivial because there are only a few entries. The rest time (around 70 ms) is spent on copying PTEs. Based on the results, we have the following observation.

**Observation 1.** *For the fork operation, the execution time dramatically grows with the instance size increasing, and the page table copy dominates the cost. For the page table copy, the overhead of copying PGDs and PUDs is trivial, while that of copying PMDs and PTEs is non-negligible.*

### 3.2 Impact of Fork Operations on Latency

We reveal the impact of the fork operation on the query latency in this subsection. Before presenting the results, we first introduce existing optimization approaches to the fork operation.

*Huge Page* [13, 14]. In the operating system (OS), we can increase the page size to reduce the number of pages used by a process, for example, setting the page size to 2MB instead of 4KB. The large page size can reduce the number of PTEs and accelerate the fork operation. However, previous study found that the page fault latency can increase from  $3.6\mu s$  to  $378\mu s$  after enabling huge page because compacting and zeroing memory in the page fault incurs higher overhead on huge page than that on regular page [30]. Moreover, Redis consumes much more memory space after enabling the huge page because applications do not always fully utilize the big page (e.g., in the experiment of [30], the memory consumption of Redis increased from 12.2GB to 20.7GB). Additionally, allocating huge page can lead to many fragments in the physical memory because it requires consecutive physical memory areas to build a huge page. Consequently, the kernel needs to perform a heavy defragment operation which leads to high CPU utilization [39]. The experiment of [39] shows that the benchmark *milc* in SPEC CPU 2006 spends 343 seconds (37% of its overall execution time) in kernel mode to perform the defragment operation when the memory is highly fragmented.

Consequently, huge page is recommended to be disabled in many databases (e.g., Couchbase, MongoDB, KeyDB and Redis [15, 27, 36, 44]). In particular, two No-SQL databases, Couchbase and MongoDB, recommend users to disable the technique because huge page performs poorly with random memory accesses in workloads [15, 36]. KeyDB and Redis, which are two IMKVSes, recommend users to disable huge page because the technique can incur a big latency penalty and big memory usage [27, 44]. Specifically, if we enable the huge page, the parent and child processes share huge pages after calling *fork* to persist on disk. In a busy instance, a few event

**Table 1: Migrating a page from “X” to “Y” when the parent and child process share the page table.**

Step	Operation	Parent(P)	Child(C)
1	Initial state	TLB: V $\rightarrow$ X PTE: V $\rightarrow$ X	TLB: V $\rightarrow$ X PTE: V $\rightarrow$ X
2	P: Set PTE $\rightarrow$ None present	TLB: V $\rightarrow$ X PTE: V $\rightarrow$ N	TLB: V $\rightarrow$ X PTE: V $\rightarrow$ N
3	P: Flush TLB	TLB: N/A PTE: V $\rightarrow$ N	TLB: V $\rightarrow$ X PTE: V $\rightarrow$ N
4	C: Skipped because N! $\neq$ X	TLB: N/A PTE: V $\rightarrow$ N	TLB: V $\rightarrow$ X PTE: V $\rightarrow$ N
5	P: Update PTE	TLB: N/A PTE: V $\rightarrow$ Y	TLB: V $\rightarrow$ X PTE: V $\rightarrow$ Y
6	P&C: Access V	TLB: V $\rightarrow$ Y PTE: V $\rightarrow$ Y	TLB: V $\rightarrow$ X PTE: V $\rightarrow$ Y

loops in either of the two processes will cause to target thousands of pages and trigger the copy operation of a large amount of process memory because of the copy-on-write mechanism in OS. Consequently, this leads to a big latency and a big memory usage. As such, our profiling does not involve this technique.

*Shared Page Table* [17, 34, 59]. This technique proposed to share the page table between parent and child in a copy-on-write (CoW) manner. Specifically, the fork operation returns immediately after copying the metadata except for the page table. The page table will be copied in a CoW manner. However, we find that the shared page table design introduces the data leakage problem, the working set size estimation problem and the NUMA problem. First, the inconsistency between the shared page table and the translation lookaside buffer (TLB) of the child process can lead to the data leakage problem. Second, we cannot accurately estimate the working set size, which indicates the memory usage of each process and is important for cloud resource management [57]. This is because the usage is calculated based on the states in the page table [20], while the table is shared by multiple processes. Third, the shared page table and the corresponding processes can locate on different NUMA nodes, which increases the TLB miss overhead [3, 38]. Moreover, the NUMA balance mechanism cannot work as expected due to the shared page table. Due to space limit, we discuss the working set size problem and the NUMA problem in the technical report [37]. In the following, we use the example in Table 1 to demonstrate the data leakage problem, which is the most serious one among the three problems. We also write a test program<sup>2</sup> to reproduce the example in practice.

TLB is the hardware to accelerate the translation from a virtual address to a physical address by caching recent translation results. Initially, the virtual address “V” is mapped to the physical address “X” in PTE and the mapping is cached in TLB. Note that PTE is shared between parent and child, whereas the two processes have their own TLB entries. Suppose that the memory management mechanism (e.g., memory compaction [5, 53], swap [55] and NUMA balance [54]) of OS mitigates the page from “X” to a new page frame “Y” in the parent process. Then, OS sets the mapping from “V $\rightarrow$ X” to “V $\rightarrow$ N” (None Present) to indicate that the mapping is invalid

<sup>2</sup><https://doi.org/10.5281/zenodo.7189585>, Last accessed on 2022/11/13.




Figure 4: The 99%-ile latencies of normal queries and snapshot queries in Redis.

in Step 2 and flushes the parent’s TLB entry in Step 3. For other processes, the OS loops over each of them to check whether its PTEs contain "V->X"; if so then set the value to "V->N" and flush the TLB; otherwise, skip the process. This works well if each process has a private page table. However, ODF uses the shared page table design. When the OS checks the child process, the PTE has been set to "V->N" in the parent process because the PTE is shared between the parent and child processes and therefore the OS cannot find any PTE with the value "V->X" in the child process. Thus, the OS skips the update to the child in Step 4. In Step 5, the parent updates PTE to map "V" to "Y". Although PTE has the correct mapping, the child’s TLB entry is inconsistent with PTE. Consequently, the future access to "V" in the child can lead to a data leakage problem. Despite that, we study the performance of the shared page table-based *fork* in our experiments for comparison purposes. As ODF [59] is the latest work and the only one that is publicly available<sup>3</sup> among the methods [17, 34, 59] adopting the shared page table design, our experiments focus on ODF.

**Experiment Results.** Figures 4 and 5 present the 99%-ile latency and the maximum latency of normal queries and snapshot queries, respectively. Snapshot-ODF denotes the results of using the On-Demand-Fork [59] to take the snapshot, which is the latest shared page table-based fork method. Snapshot-DEF denotes the results of using the default *fork* in Linux. As shown in the figures, the latency of normal queries slightly increases with the instance size varying from 1GB to 64GB. In contrast, the value of snapshot queries grows sharply. The shared page table technique dramatically reduces the latency of the default *fork*, especially for the large instance. For example, on the 64GB instance size, the optimization reduces the 99%-ile latency from 911.95ms to 3.96ms and the maximum latency from 1204.78ms to 59.28ms. The latency of Snapshot-ODF is higher than that of Normal because the CoW of the page table frequently interrupts the engine process. For example, the engine process is interrupted over 7000 times on the 16GB instance, which leads to frequent out-of-service for queries. According to the analysis of existing optimization methods and the experiment results, we have the following observation.

**Observation 2.** *The fork operation has a significant impact on the latency of snapshot queries, and the tail latency of the default function is up to hundreds of milliseconds. Although the shared page table technique reduces the latency of the default fork operation, the*

<sup>3</sup><https://github.com/rssys/on-demand-fork>, Last accessed on 2022/11/13.




Figure 5: The maximum latencies of normal queries and snapshot queries in Redis.

*overhead incurred by frequent interruptions is non-negligible and the shared page table introduces the potential data leakage problem.*

### 3.3 Summary

Based on the observations, we find that copying the page table dominates the execution time of the default *fork* in Linux, especially, for large instances. The overhead of *fork* results in a long latency (up to hundreds of milliseconds) for queries issued during the invocation of *fork*. Although several optimization methods [13, 14, 17, 34, 59] of the fork operation have been proposed, they either lead to poor performance of IMKVSes, or have a data leakage problem, which potentially generates an inconsistent snapshot. Thus, these optimizations cannot be used in the fork-based snapshot mechanism in IMKVSes. As IMKVSes have a rigid requirement on the latency to serve online scenarios, a high-performance *fork* is required to reduce the long latency of snapshot queries.

## 4 DESIGN OF ASYNC-FORK

In this section, we introduce the design of Async-fork, an operating system-based solution that effectively reduces the long latency of snapshot queries without incurring extra vulnerabilities.

### 4.1 General Idea

Figure 1(c) shows the general idea of **Async-fork**. While the parent process is responsible to copy page tables during snapshot in the default *fork*, as shown in the figure, Async-fork offloads the work of copying page table to the idle child process while keeping other steps in the fork unchanged. In this way, the parent process is able to handle queries while the child process copies the page table from the parent process simultaneously. However, it is non-trivial to achieve the design of Async-fork, as a snapshot may be inconsistent due to the asynchronous operations on the page table.

The inconsistency happens when the parent process modifies a PTE before the child process has copied it. Take Figure 6 as an example. The IMKVS takes a snapshot at time  $T_0$ , and the in-memory data is  $\{(k_0, v_0), (k_1, v_1)\}$ . Suppose a user query that sets a new KV pair  $(k_2, v_2)$  arrives at time  $T_2$  (① in Figure 6), and the child process is copying PTEs from time  $T_1$  to  $T_3$ . The parent process handles this query, and PTE<sub>2</sub> is modified to point to a new page (Page<sub>2</sub>) that contains  $(k_2, v_2)$ . If the child process has not copied the original PTE<sub>2</sub> before the modification, it would copy the modified PTE<sub>2</sub>, and owns the new pair  $(k_2, v_2)$  in its memory space. In this way, the




Figure 6: The challenges in Async-fork.

data  $(k_2, v_2)$  is persisted, and inconsistency happens (the key-value pair  $(k_2, v_2)$  does not exist when the snapshot is taken at time  $T_0$ ).

Async-fork resolves the above inconsistency problem by using the parent process to proactively synchronize the modified PTEs to the child process. We explain the detailed steps in Section 4.2. Two main challenges have to be resolved in this solution.

Firstly, it is necessary to detect all the PTE modifications. However, besides of the user operations, many inherent memory management operations in the operating system also cause PTE modifications. For instance, the OS periodically migrates pages among NUMA nodes [54], causing the involved PTEs to be modified as inaccessible (② in Figure 6). We describe the method to detect the PTE modifications in Section 4.3.

Secondly, errors may occur during Async-fork. For instance, the child process may fail to copy an entry due to out of memory (③ in Figure 6). In this case, error handling is necessary, as we should restore the process to the state before it calls Async-fork. We present how errors are handled in Section 4.4.

By resolving the above challenges, the time of the parent process used on *fork* is greatly reduced, and the latency of the snapshot queries can be reduced in consequence. Moreover, after the child process finishes copying the page table, the two processes have their complete private page table. Therefore, the Async-fork does not introduce the data leakage vulnerability.

In the following sections, we use the terminology in x86 Linux to explain our design. We also implement Async-fork on Arm Linux using similar ideas, while similar results are achieved compared with x86. Algorithm 1 describes how the parent process (Line 1 to 14) and the child process (Line 15 to 24) work in Async-fork.

## 4.2 Proactive Synchronization

Before introducing the proactive synchronization in detail, we first introduce how we offload the work of copying page table to the child process in Async-fork.

**Copying Page Table Asynchronously.** In the default *fork*, the parent process traverses all its VMAs and copies the corresponding parts of the page table to the child process. The page table is copied from top to bottom. In Async-fork, the parent process roughly

### Algorithm 1: The framework of Async-fork

```

// The parent process
1 for each VMAi in (VMA0, VMA1, VMA2...) do
2   Copy VMAi to child process;
3   Copy VMAi's PGDs/PUDs to child process;
4   Set all VMAi's PMDs to be write-protected;
5 end
6 Put the child process on a CPU to run it;
7 while true do
8   if PTE modification is detected and this PTE's PMD is write-protected then
9     Copy PMD to child process;
10    Copy 512 PTEs of this PMD to child process;
11    Set the PMD to be writeable;
12    Set the 512 PTEs to be write-protected;
13  end
14 end
// The child process
15 for each VMAi in (VMA0, VMA1, VMA2...) do
16   for each PMDj of VMAi do
17    if PMDj is write-protected then
18      Copy PMDj from parent process;
19      Copy 512 PTEs of PMDj from parent process;
20      Set the PMD to be writeable;
21      Set the 512 PTEs to be write-protected;
22    end
23  end
24 end

```

follows the above process, but only copies PGD, P4D (if exists) and PUD entries to the child process (Lines 1 to 3 in Algorithm 1). After that, the child process starts to run, and the parent process returns to the user mode to handle queries (Lines 7 to 14). The child process then traverses the VMAs, and copies PMD entries and PTEs from the parent process. Figure 7 shows an example of the asynchronous page table copy. In the figure, the PGD/PUD entries have been copied by the parent process, but PMD entries and PTEs have not yet (some PMD entries point to "null"). The child process copies PMDs/PTEs from the parent process (e.g., PMD<sub>0</sub> and its PTEs).

We offload the work of copying PMD entries and PTEs to the child process because the overhead of copying them is non-negligible, as analyzed in Section 3.1. Meanwhile, we keep that the parent process copies PGD/PUD entries. This is because the overhead of copying PGD/PUD entries is trivial, and it is more robust to minimize the change to the Linux kernel.

**Synchronizing Modified PTEs Proactively.** When the child process is responsible for copying PMD entries and PTEs, it is possible that the PTEs are modified by the parent process before they are actually copied. Note that only the parent process is aware of the modifications (the way to detect the PTE modification will be introduced in Section 4.3).

In general, there are two ways to copy the to-be-modified PTEs for the consistency. 1) the parent process proactively copies the PTEs to the child process; 2) the parent process notifies the child process to copy the PTEs and waits until the copying is finished. As both ways result in the same interruption in the parent process, we choose the former way (Lines 8 to 12 in Algorithm 1).

More specifically, when a PTE is modified during snapshot, the parent process copies not only this PTE but also all the other PTEs of a same PTE table (512 PTEs in total), as well as the parent PMD entry to the child process proactively. For instance, when PTE<sub>1</sub> in




Figure 7: An example of copying page table in Async-fork. “RW=1” represents writable and “RW=0” represents write-protected.

Figure 7 is modified, the parent process proactively copies PMD<sub>2</sub>, PTE<sub>0</sub> and PTE<sub>1</sub> to the child process. We choose to copy the entire PTE table because we can quickly detect a range of PTEs that will be modified, but accurately identifying which one will be modified is expensive in practice.

**Eliminating Unnecessary Synchronizations.** Always letting the parent process copy the modified PTEs is unnecessary, as it is possible that the to-be-modified PTEs have already been copied by the child process. We identify if the PMD entries and PTEs have been copied by the child process, to avoid unnecessary synchronizations. A flag is required to track this status. We reuse the R/W flag of the PMD entry to record the status. Since the R/W flag is only used when the PMD entry points to a huge page in the x86 Linux kernel, the flag tracks the status correctly. We reuse the R/W flag because 1) this design can avoid adding new fields to the kernel data structures; and 2) the popular databases (e.g., Redis, KeyDB, MongoDB, and Couchbase) recommend disabling the huge page to improve the performance [15, 27, 36, 44]. Additionally, it is unnecessary to use Async-fork if applications use the huge page because the applications with huge page do not require PTEs but a small number of PGD/PUD/PMD entries (i.e., the page table is small).

An alternative approach is to use an unused bit in the *struct page* as the flag. Specifically, each PMD entry points to a PTE table, while OS maintains a data structure (*struct page*) for each PTE table. The *struct page* has bits that are not used by the current Linux kernel. Previous research uses these bits as flags, for example, ODF uses some bits in *struct page* as a reference counter. However, this approach requires further modification to the kernel to initialize the bit. Therefore, we do not adopt the design using the *struct page*.

If a PMD entry and its 512 PTEs have not been copied to the child process, the PMD entry will be set as write-protected (e.g.,

PMD<sub>1</sub> in Figure 7). Note that, it does not break the CoW strategy of *fork* since it still triggers the page fault when the corresponding page is written on x86 [21]. Once the PMD/PTEs have been copied to the child process (e.g., PMD<sub>0</sub>), the PMD entry is changed to be writable (the PTEs are changed to be write-protected to maintain the CoW strategy). Since both parent and child processes lock the page of the PTE table with *trylock\_page()* when they are copying PMD entries and PTEs, they will not copy PTEs pointed by the same PMD entry at the same time.

### 4.3 Detecting Modified PTEs

The operations that modify PTEs in the OS can be divided into two categories: 1) *VMA-wide modification*. Some operations act on specific VMAs, including creating, merging, deleting VMAs and so on. The modification of a VMA may also cause the VMA’s PTEs be modified. For example, the user sends queries to delete lots of KV pairs. The IMKVS (parent process) then reclaims the corresponding virtual memory space by *munmap*. Some VMAs are hence split or deleted while the VMAs’ PTEs are deleted as well. A VMA is usually large because the operating system always tends to merge adjacent VMAs. It means that the VMA-wide modification usually causes extensive PTE modifications, while many VMA’s PMD entries are involved. 2) *PMD-wide modification*. Other operations modify the PTE directly. For example, the page of parent process can be reclaimed by the out of memory (OOM) killer. In this case, one PMD entry is involved. Note that swapping or migrating a 4KB page will change the PTE but the data will not be changed, so we will not handle it. Due to limited space, we summarize the locations where operations in the OS modify VMAs/PTEs as checkpoints in the technical report [37].

We implement the detection by hooking the checkpoints. Once a checkpoint is reached, the parent process checks whether the involved PMD entries and PTEs have been copied (by checking the R/W flag of the PMD entry). For a VMA-wide modification, all the PMD entries of this VMA are checked, while only one PMD entry is checked for a PMD-wide modification. The uncopied PMD entries and PTEs will be copied to the child process before modifying them.

If a VMA is large, the parent process may take a relatively long time to check all PMD/PTE entries by looping over each of them. We therefore introduce a two-way pointer, which helps the parent process quickly determine whether all entries of a VMA have been copied to the child, to reduce the cost. Each VMA has a two-way pointer, which is initialized by the parent process during the invocation of the Async-fork function. The pointer in the VMA of the parent process (resp. the child process) points to the corresponding VMA of the child process (resp. the parent process). In this way, the two-way pointer maintains a connection between the VMAs of the parent and the child. The connection will be closed after all PMDs/PTEs of the VMA are copied to the child. Specifically, if no VMA-wide modification happens during the copy of PMDs/PTEs of the VMA, then the child closes the connection by setting the pointers in the VMAs of both the parent and child to *null* after the copy operation. Otherwise, the parent will synchronize the modification (i.e., copying the uncopied PMDs/PTEs to the child), and close the connection by setting the pointers to *null* after the copy operation. As both parent and child processes can access the



two-way pointers, the pointers are protected by locks to keep the state consistent. When a VMA-wide modification occurs, the parent process can quickly determine whether all PMDs/PTEs of a VMA have been copied to the child by checking the pointer's value, instead of looping over all these PMDs. Besides, the pointer is also used in handling errors (see Section 4.4).

#### 4.4 Handling Errors

Since copying the page table involves memory allocation, some errors may occur during both the default *fork* and Async-fork. For instance, a process may fail to initialize a new PTE table due to out of memory. Such error may only happen in the parent process in the default *fork* and has a standard way to handle the error. However, the copying of page table is offloaded to the child process in Async-fork, such error may happen in the child process, and a method is required to handle such errors. Specifically, we should restore the parent process to the state before it calls Async-fork, to ensure that the parent process will not crash in the future.

As Async-fork may modify the R/W flags of the PMD entries of the parent process, we roll back these entries to be writable when errors occur in Async-fork. Errors may occur 1) when the parent process copies PGD/PUD entries, 2) when the child process copies PMD/PTEs, and 3) during a proactive synchronization. In the first case, the parent process rolls back all the write-protected PMD entries. In the second case, the child process rolls back all the remaining uncopied PMD entries. After that, we send a signal (*SIGKILL*) to the child process. The child process will be killed when it returns to the user mode and receive the signal. In the third case, the parent process only rolls back the PMD entries of the VMA containing the PMD entry that is being copied. The purpose is to avoid contending for the PMD entry lock with the child process. An error code is then stored into the two-way pointer of the VMA. Before (and after) copying PMDs/PTEs of a VMA, the child process will check the pointer to see whether there are errors. If so, then it stops copying PMD/PTE entries immediately and performs the rollback operations that are already described in the second case.

### 5 OPTIMIZATION AND IMPLEMENTATION

In this section, we present the optimization that further improves the performance of Async-fork, and the way to implement Async-fork in Linux.

#### 5.1 Accelerating Page Table Copy

The parent process is still interrupted when a proactive PTE synchronization is triggered. A straightforward way to reduce the cost of the proactive PTE synchronization is to let the child process to first copy the PTEs potentially modified before other PTEs. However, this method is not practical because the data accessed by user queries are relatively random and the child process cannot determine which PTEs will be modified when the parent invokes Async-fork. We therefore propose to minimize the number of proactive PTE synchronizations by reducing the duration of copying the page table.

As VMAs are independent, the kernel threads can totally perform the copy in parallel and obtain near-linear speedup. Therefore, the child process may launch multiple kernel threads to copy




Figure 8: The implementation of Async-fork.

PMD/PTEs in multiple VMAs in parallel, so that the copy completes faster. It can effectively reduce the number of proactive PTE synchronization because the synchronization only happens during the period that the child process is copying PMD/PTEs. The probability of triggering a proactive PTE synchronization gets lower when the period becomes shorter. The experiment in Section 6.3 shows the efficiency of this optimization.

Multiple kernel threads consume CPU cycles. These threads periodically check whether they should be preempted and give up CPU resources by calling `cond_resched()`, in order to reduce the interference on other normal processes.

#### 5.2 Implementation of Async-fork

Figure 8 shows the implementation of Async-fork with a modular design. We encapsulate the code ingested to the kernel into a hook function, which is instantiated in a kernel module. We can insert/remove the module in/from the kernel as necessary. The hook function is inserted to the call path of default *fork* and memory subsystem in the Linux kernel (version 4.19). Users can determine the usage of the default *fork* and Async-fork with the parameter *F*.

The hook function is enhanced from `copy_page_range()`, which accepts a flag (Fast or Slow) to control the copying of the page table.

With “Fast” flag, the function copies PGD/PUD entries and sets PMD entries to be write-protected. With “Slow” flag, the function copies the write-protected PMD entries and PTEs; when the copying finishes, it sets the PMD entries to be writable and sets the PTEs to be write-protected. When Async-fork is called, the parent process copies the page table to the child process using the function with “Fast” flag. At the end of Async-fork’s invocation, the parent process puts the child process into the *runqueue* of a CPU and returns to the user mode. Before the child process returns to the user mode, it copies the page table using the function with “Slow” flag. The parent process proactively copies the uncopied PMD/PTEs (using the function with “Fast” flag) to the child process before modifying them.

**Flexibility.** For the IMKVS workload that has small memory footprint, the page table copy is already short. In this case, Async-fork brings small benefit. For these workloads, we provide an interface in *memory cgroup* to control whether Async-fork is enabled

(as well as the number of kernel threads used to speed up copying PMD/PTEs in the child process). Specifically, when users add a process to a *memory cgroup*, they can pass a parameter to enable/disable Async-fork at run time (i.e., the parameter  $F$  in Figure 8). As shown in the figure, if the parameter value is 0, then the process will use the default *fork*. Otherwise, Async-fork is enabled. As such, users can determine which fork operation the process uses as necessary and use Async-fork without any modification in the source code of applications. The process will use the default *fork* if no parameter is passed in.

**Memory overhead.** The only memory overhead of Async-fork comes from the added pointer (8B) in each VMA of a process. In the case, the memory overhead in a process is the number of VMAs times 8B. Considering a machine with 512GB main memory while 400 processes run simultaneously, there are roughly 760,000 VMAs according to our statistics. In this case, the memory overhead will be  $760000 \times 8B \approx 6MB$ . This overhead is generally negligible.

**Support for ARM64.** The design of Async-fork can also be implemented on ARM64. Specifically, we use the  $APT_{table}[1:0]$  [33] in the table descriptor of the PMD entry as the R/W flag. Async-fork can also be implemented on other architectures that support hierarchical attributes in the page table.

**Consecutive Snapshots.** It is possible that the parent process starts the next snapshot using Async-fork before the previous child process finishes copying PMD/PTEs, as the parent process returns to the user mode before the child process. In the current implementation, we will not block the next Async-fork call but keep a VMA’s page table be copied by only one child process at any time. When the parent process copies a VMA to the child process it checks the two-way pointer to identify whether there exists a previous child process copying the page table of this VMA. If exists, it proactively copies the whole page table of this VMA to the previous child process. Async-fork adopts this design because supporting concurrent fork operations need to track all child processes that are copying page tables. This is hard because the kernel does not provide the information and we need to inject new data structures into the kernel to record the states and synchronize these child processes.

As Async-fork cannot support concurrent fork operations in a process (the same as the default *fork* in the OS), Async-fork cannot support the cases where the parent process needs to conduct the fork operation in an ultra-high frequency (e.g., in milliseconds). In spite of the limitation, Async-fork works well for database use cases because a storage engine does not create many child processes from the parent process simultaneously in practice. Specifically, conducting the fork operation frequently leads the parent to frequently turn into the kernel mode, which hurts the service quality. Moreover, many processes executing in parallel will degrade the performance due to the resource contention (e.g., IO bandwidth and CPU resources). Therefore, IMKVSes (e.g., Redis and KeyDB) do not recommend ultra-frequent data snapshots (generally 60 seconds [47]) and have no cases requiring the parent process to invoke Async-fork in milliseconds to our knowledge. For HyPer, which uses the fork operation to support concurrent transaction processing, Async-fork can work well. This is because the parent process handles OLTP (a single updater) that has a rigid requirement on the latency, whereas the child process executes OLAP which generally has a long execution time and is more tolerant of the latency than

OLTP. Moreover, HyPer notices the cost of the fork operation and designs a novel mechanism, which makes multiple OLAP queries to share a snapshot, to improve the performance.

## 6 EVALUATION

In this section, we evaluate the effectiveness of Async-fork. We conduct experiments on our local machine for the purpose of test flexibility.

### 6.1 Experimental Setup

We evaluate Async-fork on a machine with two Intel Xeon Platinum 8163 processors, each of them has 24 physical cores (48 logical cores). The machine has 384GB memory, and 1TB NVMe hard drive. In terms of software setup, the experimental platform runs CentOS 7.9 with Linux 4.19.

**Benchmarks.** We use Redis (version 5.0.10) [42] and KeyDB (version 6.2.0) [24] compiled with gcc 6.5.1 as the representative IMKVS servers, and use Redis benchmark [41] as well as Memtier benchmark [46] to be the workload generators. The benchmarks reveal the scenario where multiple clients send requests to the IMKVS server simultaneously. Similar to prior work [12], we enhance Redis benchmark to generate queries in an open-loop mode for measuring the latency accurately [48, 58].

By default, the experiments are conducted with the following settings: 1) 50 clients (default settings) are used in Redis benchmark, while 50 clients are used in Memtier benchmark for consistency. 2) The key range is set to  $2 \times 10^8$ , the key size of 8B and the value size of 1024B. 3) Each experiment is repeated by five times and the average results are reported. 4) In Async-fork, the child process launches 7 additional kernel threads (together with the child process itself, there are 8 threads in total) to help it copy PMDs/PTEs faster. Our experiment in Figure 14 shows that Async-fork still outperforms the state-of-the-art solution with a single child process. 5) The KeyDB server is configured with 4 threads.

**Metrics.** We launch a large number of queries ( $5 \times 10^6$ ) to the IMKVS and record the latency of the queries that arrive in the snapshot process (start from the parent process calls *fork* until the child process persists all in-memory data). Specifically, we measure the 99%-ile latency because latency-sensitive services generally provide an SLA on some percentile. Moreover, we report the maximum latency because IMKVS is often used for demanding use cases that have a rigid requirement for the worst-case latency [43]. For example, the increase of the maximum latency of Redis can lead to read error on connections [52]. Therefore, in the production environment [19], the maximum latency is an important indicator of system stability. Due to limited space, we report the experiment results of query processing throughput and total out-of-service time of the parent process in the technical report [37]. In the report, we also discuss the method of tuning IMKVSes to further improve the performance.

**Baselines.** We compare Async-fork with On-Demand-Fork (denoted by *ODF* in short) [59], the state-of-the-art shared page table-based fork. In ODF, each time a shared PTE is modified by a process, not only one PTE but 512 PTEs located on the same PTE table will be copied at the same time. We do not report the results of the default *fork* in this section since it results in 10X higher latency




Figure 9: The 99%-ile latency of snapshot queries.




Figure 10: The maximum latency of snapshot queries.

compared with both Async-fork and ODF in most cases (already presented in Section 3.2). To accurately measure the performance of the IMKVSes in the experiments, we trigger the snapshot operation manually using the BSAVE command.

## 6.2 Overall Evaluation

We first evaluate Async-fork using write-intensive workloads that require frequent snapshots for data persistence.

**Latency Results.** In this experiment, we use the Redis benchmark [41] to generate the write-intensive workload by issuing *SET* queries to the IMKVSes, and configure the clients to send 50,000 such queries in a second. Since the IMKVS is often used in the context of demanding use cases, there are usually strict requirements on both the 99%-ile latency and the worst case latency [25, 43].

Figure 9 shows the 99%-ile latencies of snapshot queries in Redis and KeyDB, with ODF and Async-fork. As observed, Async-fork outperforms ODF in all the cases, and the performance gap increases when the instance size gets larger. For instance, operating on a 64GB IMKVS instance, the 99%-ile latency of the snapshot queries is 3.96ms (Redis) and 3.24ms (KeyDB) with ODF, while the 99%-ile latency reduces to 1.5ms (Redis, 61.9% reduction) and 1.03ms (KeyDB, 68.3% reduction) with Async-fork.

Figure 10 shows the maximum latency of the snapshot queries. As observed, Async-fork greatly reduces the maximum latency of the benchmarks compared with ODF, even if the instance size is small. For a 1GB IMKVS instance, the maximum latencies of the snapshot queries are 13.93ms (Redis) and 10.24ms (KeyDB) respectively with ODF, while the maximum latencies are decreased to 5.43ms (Redis, 60.97% reduction) and 5.64ms (KeyDB, 44.95% reduction) with Async-fork.

**Deep Diving.** The interruption of Async-fork is caused by the proactive synchronization, whereas the interruption of ODF is caused by the CoW of the page table. We measure the interruptions of the parent process during the snapshot to understand the reason that Async-fork outperforms ODF. Specifically, the parent process turns into kernel mode and is out of service for queries when executing the *copy\_pmd\_range()* function. The overhead of *copy\_pmd\_range()* dominates the cost of executing in the kernel mode. In order to examine the number of interruptions and the




Figure 11: The frequency of interruptions in the parent process during the snapshot process.




Figure 12: The 99%-ile latency and maximum latency of snapshot queries under different workloads in an 8GB IMKVS.

out-of-service time, we use the *bcc* tool [2] to count the number of *copy\_pmd\_range()* invocations and measure the execution time of each invocation. The result of *bcc* is a histogram in which the bucket is the time duration and the frequency is the number of invocations whose execution time falls into the bucket. The categories [16us, 31us] and [32us, 63us] are two default buckets of *bcc*. In our experiments, all invocations fall into the two buckets.

Figure 11 shows the frequency of the interruptions within [16us, 31us] and [32us, 63us]. We can see that Async-fork significantly reduces the frequency of interruptions. For example, Async-fork reduces the frequency of interruptions from 7348 to 446 on the 16GB instance. Async-fork greatly reduces the interruptions because the interruptions happen only when the child process is copying PMD/PTEs (the required time is within 600ms as in Figure 15). However, the interruption can happen until all data is persisted by the child process in ODF, while the data persistence operation requires tens of seconds (e.g., persisting 8GB in-memory data takes about 40s). Under the same workload, the parent process is more vulnerable to interruption when using ODF.

## 6.3 Detailed Evaluation

**Sensitivities to the Read/Write Patterns.** Figure 12 shows the results using four workloads with different read-write patterns generated with Memtier [46]. In the figure, “1:1 (Uni.)” represents the workload with 1:1 Set:Get Ratio and the uniform random access pattern, while “1:10 (Gau.)” is the workload with 1:10 Set:Get ratio and the Gaussian distribution access pattern.

Observed from Figure 12, Async-fork still outperforms ODF. The benefit is smaller for the workload with more *GET* queries. This




Figure 13: The 99%-ile latency and the maximum latency of snapshot queries under different numbers of clients in an 8GB IMKVS.




Figure 14: The 99%-ile and maximum latency of snapshot queries when Async-fork uses 1 or 8 threads.

is because the serving (parent) process only copy a small number of PTEs for the *GET*-intensive workloads. Moreover, the modified memory is smaller in the experiment with the Gaussian distribution compared with the uniform random access pattern. With random pattern, the key-value pairs in the IMKVS have the same probability of being accessed, and parts of key-value pairs may be accessed repeatedly with the Gaussian Distribution access pattern. Since shared PTEs are copied only when they are modified for the first time in ODF, the parent process is interrupted fewer times with the Gaussian distribution access pattern.

In general, Async-fork works better for write-intensive workloads. The larger the modified memory is, the better Async-fork performs. Integrating Async-fork with CCoW [22] to improve the performance on write-intensive workloads is an interesting research direction because Async-fork can utilize CCoW to copy the PTEs of high-locality memory pages in advance to reduce the number of proactive synchronizations.

**The Impact of the Number of Clients.** In this experiment, we change the number of clients in Redis-benchmark while keeping sending 50,000 *SET* queries every second to an 8GB IMKVS server. Figure 13 shows the results of 99%-ile and maximum latency with 10, 50, 100 and 500 clients. As observed, Async-fork outperforms ODF, while the performance gap increases as the number of clients increases. This is because more requests arrive at the IMKVS at the same time when the number of clients increases. As a result, more PTEs may be modified at the same time, and the duration of one interruption to the parent process may become longer.




Figure 15: (a) The time that the child process takes to copy PMDs/PTEs in Async-fork. (b) The 99%-ile and maximum latency in an 8GB Redis instance.

**The Impact of the Number of Threads In the Child Process.** In Async-fork, multiple kernel threads may be used to copy PMD/PTEs in parallel in the child process. Figure 14 shows the 99%-ile and maximum latency of snapshot queries under different Redis instance sizes. Async-fork#*i* represents the results of using *i* threads in total to copy the PMDs/PTEs.

Observed from Figure 14, Async-fork#1 (the child process itself) still brings shorter latency than ODF. The maximum latency of the snapshot queries is decreased by 34.3% on average, compared with ODF. We can also find that using more threads (Async-fork#8) can further decrease the maximum latency. This is because the sooner the child process finishes copying PMDs/PTEs, the lower probability the parent process is interrupted to proactively synchronize PTEs.

In more detail, Figure 15(a) shows the time that the child process takes to copy PMDs/PTEs with different numbers of kernel threads, while Figure 15(b) shows the corresponding 99%-ile latency and maximum latency in an 8GB Redis instance. As we can see, launching more kernel threads effectively reduces the time of copying PMDs/PTEs in the child process. The shorter the time is, the lower the latency becomes.

## 7 CONCLUSION

In this paper, we study the latency spikes incurred by the fork-based snapshot mechanism in IMKVSes and address the problem from the operating system level. In particular, we conduct an in-depth study to reveal the impact of the fork operation on the latency spikes. According to the study, we propose Async-fork. It optimizes the fork operation by offloading the workload of copying the page table from the parent process to the child process. To guarantee data consistency between the parent and the child, we design the proactive synchronization strategy. Async-fork is implemented in the Linux kernel and deployed in production environments. Extensive experiment results show that Async-fork can significantly reduce the tail latency of queries arriving during the snapshot period.

## ACKNOWLEDGMENTS

This work was partially sponsored by the National Natural Science Foundation of China (62232011, 62022057), and Shanghai international science and technology collaboration project (21510713600). This work was also supported by Alibaba Group through Alibaba Innovative Research Program. Pu Pang and Gang Deng contributed equally to this work. Quan Chen and Minyi Guo are the corresponding authors.

## REFERENCES

- [1] [n.d.]. fork(2). <https://linux.die.net/man/2/fork>.
- [2] 2022. BPF Compiler Collection (BCC). <https://github.com/iovisor/bcc>.
- [3] Reto Achermann, Ashish Panwar, Abhishek Bhattacharjee, Timothy Roscoe, and Jayneel Gandhi. 2020. Mitosis: Transparently self-replicating page-tables for large-memory machines. In *Proceedings of the 25th International Conference on Architectural Support for Programming Languages and Operating Systems*. 283–300.
- [4] Antirez. [n.d.]. Redis persistence demystified. <http://antirez.com/post/redis-persistence-demystified.html>.
- [5] Vlastimil Babka. 2016. mm, compaction: introduce kcompactd. <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=698b1b30642f1ff0ea10ef1de9745ab633031377>.
- [6] Katelin A Bailey, Peter Hornyack, Luis Ceze, Steven D Gribble, and Henry M Levy. 2013. Exploring storage class memory with key value stores. In *Proceedings of the 1st Workshop on Interactions of NVM/FLASH with Operating Systems and Workloads*. 1–8.
- [7] Philip A Bernstein, Vassos Hadzilacos, and Nathan Goodman. 1987. *Concurrency control and recovery in database systems*. Vol. 370. Addison-wesley Reading.
- [8] Greg Bronevetsky, Rohit Fernandes, Daniel Marques, Keshav Pingali, and Paul Stodghill. 2006. Recent advances in checkpoint/recovery systems. In *Proceedings 20th IEEE International Parallel & Distributed Processing Symposium*. IEEE, 8–pp.
- [9] Tuan Cao. 2013. *Fault tolerance for main-memory applications in the cloud*. Cornell University.
- [10] Tuan Cao, Marcos Vaz Salles, Benjamin Sowell, Yao Yue, Alan Demers, Johannes Gehrke, and Walker White. 2011. Fast checkpoint recovery algorithms for frequently consistent applications. In *Proceedings of the 2011 ACM International Conference on Management of Data*. 265–276.
- [11] Badrish Chandramouli, Guna Prasaad, Donald Kossmann, Justin Levandoski, James Hunter, and Mike Barnett. 2018. Faster: A concurrent key-value store with in-place updates. In *Proceedings of the 2018 ACM International Conference on Management of Data*. 275–290.
- [12] Shuang Chen, Christina Delimitrou, and José F Martínez. 2019. Parties: Qos-aware resource partitioning for multiple interactive services. In *Proceedings of the 24th International Conference on Architectural Support for Programming Languages and Operating Systems*. 107–120.
- [13] The Kernel Development Community. 2021. HugeTLB Pages. <https://www.kernel.org/doc/html/latest/admin-guide/mm/hugetlbpage.html>.
- [14] The Kernel Development Community. 2021. Transparent Hugepage Support. <https://www.kernel.org/doc/html/latest/admin-guide/mm/transhuge.html>.
- [15] Couchbase. 2021. Disabling Transparent Huge Pages (THP). <https://docs.couchbase.com/server/current/install/thp-disable.html>.
- [16] Cristian Diaconu, Craig Freedman, Erik Ismert, Per-Ake Larson, Pravin Mittal, Ryan Stonecipher, Nitin Verma, and Mike Zwilling. 2013. Hekaton: SQL server’s memory-optimized OLTP engine. In *Proceedings of the 2013 ACM International Conference on Management of Data*. 1243–1254.
- [17] Xiaowan Dong, Sandhya Dwarkadas, and Alan L Cox. 2016. Shared address translation revisited. In *Proceedings of the 11th European Conference on Computer Systems*. 1–15.
- [18] Bin Fan, David G Andersen, and Michael Kaminsky. 2013. Memc3: Compact and concurrent memcache with dumber caching and smarter hashing. In *10th USENIX Symposium on Networked Systems Design and Implementation*. 371–384.
- [19] Caixin Gong, Chengjin Tian, Zhengheng Wang, Sheng Wang, Xiyu Wang, Qiulei Fu, Wu Qin, Long Qian, Rui Chen, Jiang Qi, Ruo Wang, Guoyun Zhu, Chenghu Yang, Wei Zhang, and Feifei Li. 2022. Tair-PMem: a fully durable non-volatile memory database. *Proceedings of the VLDB Endowment* 15, 12 (2022), 3346–3358.
- [20] Brendan Gregg. 2018. Working Set Size Estimation. <https://www.brendangregg.com/wss.html>.
- [21] Part Guide. 2011. Intel® 64 and ia-32 architectures software developer’s manual. *Volume 3B: System programming Guide, Part 2*, 11 (2011).
- [22] Minjong Ha and Sang-Hoon Kim. 2022. CCoW: Optimizing Copy-on-Write Considering the Spatial Locality in Workloads. *Electronics* 11, 3 (2022).
- [23] Alan Harris. 2010. Distributed caching via memcached. In *Pro ASP. NET 4 CMS*. Springer, 165–196.
- [24] Snap Inc. 2022. KeyDB. <https://keydb.dev/>.
- [25] Snap Inc. 2022. Latency Monitoring Tool. <https://docs.keydb.dev/docs/latency-monitor/>.
- [26] Snap Inc. 2022. Persistence. <https://docs.keydb.dev/docs/persistence>.
- [27] Snap Inc. 2022. Troubleshooting Latency Issues. <https://docs.keydb.dev/docs/latency/>.
- [28] Harshad Kasture and Daniel Sanchez. 2016. Tailbench: a benchmark suite and evaluation methodology for latency-critical applications. In *2016 IEEE International Symposium on Workload Characterization*. 1–10.
- [29] Alfons Kemper and Thomas Neumann. 2011. HyPer: A hybrid OLTP&OLAP main memory database system based on virtual memory snapshots. In *2011 IEEE 27th International Conference on Data Engineering*. IEEE, 195–206.
- [30] Youngjin Kwon, Hangchen Yu, Simon Peter, Christopher J Rossbach, and Emmett Witchel. 2016. Coordinated and efficient huge page management with ingens. In *12th USENIX Symposium on Operating Systems Design and Implementation*. 705–721.
- [31] Liang Li, Guoren Wang, Gang Wu, and Ye Yuan. 2018. Consistent snapshot algorithms for in-memory database systems: Experiments and analysis. In *2018 IEEE 34th International Conference on Data Engineering*. IEEE, 1284–1287.
- [32] A-P Lienes and Antoni Wolski. 2006. Siren: A memory-conserving, snapshot-consistent checkpoint algorithm for in-memory databases. In *2006 IEEE 22nd International Conference on Data Engineering*. IEEE, 99–99.
- [33] Arm Limited. [n.d.]. Arm Architecture Reference Manual Armv8, for A-profile architecture. <https://developer.arm.com/documentation/ddi0487/gb/>.
- [34] Dave McCracken. 2003. Sharing page tables in the linux kernel. In *Linux Symposium*. 315.
- [35] Michael J Mior and Eyal de Lara. 2011. Flurrydb: a dynamically scalable relational database with virtual machine cloning. In *Proceedings of the 4th Annual International Conference on Systems and Storage*. 1–9.
- [36] MongodB. 2021. Disable Transparent Huge Pages (THP) – MongoDB Manual. <https://docs.mongodb.com/manual/tutorial/transparent-huge-pages>.
- [37] Pu Pang, Gang Deng, Kaihao Bai, Quan Chen, Shixuan Sun, Bo Liu, Yu Xu, Hongbo Yao, Zhengheng Wang, Xiyu Wang, Zheng Liu, Zhuo Song, Yong Yang, Tao Ma, and Minyi Guo. 2023. Async-fork: Mitigating Query Latency Spikes Incurred by the Fork-based Snapshot Mechanism from the OS Level. [arXiv:2301.05861](https://arxiv.org/abs/2301.05861) [cs.DB]
- [38] Ashish Panwar, Reto Achermann, Arkaprava Basu, Abhishek Bhattacharjee, K Gopinath, and Jayneel Gandhi. 2021. Fast local page-tables for virtualized NUMA servers with vMitosis. In *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*. 194–210.
- [39] Ashish Panwar, Aravinda Prasad, and K Gopinath. 2018. Making huge pages actually useful. In *Proceedings of the 23th International Conference on Architectural Support for Programming Languages and Operating Systems*. 679–692.
- [40] Jiwoong Park, Yunjae Lee, Heon Young Yeom, and Yongseok Son. 2020. Memory efficient fork-based checkpointing mechanism for in-memory database systems. In *Proceedings of the 35th Annual ACM Symposium on Applied Computing*. 420–427.
- [41] RedisLab. 2021. How fast is Redis? <https://redis.io/topics/benchmarks>.
- [42] RedisLab. 2021. Redis. <https://redis.io/>.
- [43] RedisLab. 2021. Redis latency monitoring framework. <https://redis.io/topics/latency-monitor>.
- [44] RedisLab. 2021. Redis Latency Problems Troubleshooting. <https://redis.io/topics/latency>.
- [45] RedisLab. 2021. Redis Persistence. <https://redis.io/topics/persistence>.
- [46] RedisLab. 2022. memtier. [https://github.com/RedisLabs/memtier\\_benchmark](https://github.com/RedisLabs/memtier_benchmark).
- [47] RedisLab. 2022. redis.conf. <https://github.com/redis/redis/blob/7.0/redis.conf>.
- [48] Bianca Schroeder, Adam Wierman, and Mor Harchol-Balter. 2006. Open versus closed: A cautionary tale. In *3rd USENIX Symposium on Networked Systems Design and Implementation*. USENIX.
- [49] Matthew Shafer. 2012. Memcached. <https://github.com/memcached/memcached>.
- [50] Ankur Sharma, Felix Martin Schuhknecht, and Jens Dittrich. 2018. Accelerating Analytical Processing in MVCC Using Fine-Granular High-Frequency Virtual Snapshotting. In *Proceedings of the 2018 ACM International Conference on Management of Data*. 245–258.
- [51] solid IT. 2022. DB-Engines Ranking of Key-value Stores. <https://db-engines.com/en/ranking/key-value+store>.
- [52] Trivago Technology. 2017. Learn Redis the hard way (in production). <https://tech.trivago.com/2017/01/25/learn-redis-the-hard-way-in-production/>.
- [53] Linus Torvalds. 2018. linux/mm/compaction.c. <https://github.com/torvalds/linux/blob/v4.19/mm/compaction.c>.
- [54] Linus Torvalds. 2018. linux/mm/memory.c. <https://github.com/torvalds/linux/blob/v4.19/mm/memory.c>.
- [55] Linus Torvalds. 2018. linux/mm/vmscan.c. <https://github.com/torvalds/linux/blob/v4.19/mm/vmscan.c>.
- [56] Hao Zhang, Gang Chen, Beng Chin Ooi, Kian-Lee Tan, and Meihui Zhang. 2015. In-memory big data management and processing: A survey. *IEEE Transactions on Knowledge and Data Engineering* 27, 7 (2015), 1920–1948.
- [57] Wei Zhang, Ningxin Zheng, Quan Chen, Yong Yang, Zhuo Song, Tao Ma, Jingwen Leng, and Minyi Guo. 2020. URSA: Precise Capacity Planning and Fair Scheduling Based on Low-Level Statistics for Public Clouds. In *49th International Conference on Parallel Processing*. 1–11.
- [58] Yunqi Zhang, David Meisner, Jason Mars, and Lingjia Tang. 2016. Treadmill: Attributing the source of tail latency through precise load testing and statistical inference. In *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture*. IEEE, 456–468.
- [59] Kaiyang Zhao, Sishuai Gong, and Pedro Fonseca. 2021. On-demand-fork: a microsecond fork for memory-intensive and latency-sensitive applications. In *Proceedings of the 16th European Conference on Computer Systems*. 540–555.