A new approach to database management in solid-state drives

August 20 2020

Credit: Anete Lusina from Pexels

The ever-increasing workload of data centers calls for new ways to store and access data. Researchers from the Daegu Gyeongbuk Institute of Science and Technology, Korea, have developed a new approach to

manage databases in solid state drives, providing marked performance improvements in read/write delays and offloading database computation tasks from CPUs to increase efficiency and reduce power consumption.

As Web services, cloud storage and big-data services continue expanding, the gigantic hardware infrastructures they rely on—known as data centers—need to be improved to keep up with the current demand. One promising solution for improving the performance and reducing the energy load associated with reading and writing large amounts of data is to confer storage devices with some computational capabilities and offload part of the data read/write process from CPUs.

In a recent study presented at the 2020 USENIX Annual Technical Conference, researchers from Daegu Gyeongbuk Institute of Science and Technology (DGIST), Korea, describe a new way of implementing a key–value store in solid state drives (SSDs), which offers many advantages over a more widely used method.

A key–value store (also known as key–value database) is a way of storing, managing, and retrieving data in the form of key–value pairs. The most common way to implement one is through the use of a hash function, an algorithm that can quickly match a given key with its associated stored data to achieve fast read/write access.

One of the main problems of implementing a hash-based key–value store is that the random nature of the hash function occasionally leads to long delays (latency) in read/write operations. To solve this problem, the researchers from DGIST implemented a different paradigm, called "log-structured merge-tree (LSM)." This approach relies on ordering the data hierarchically, therefore putting an upper bound on the maximum latency.

In their implementation, nicknamed "PinK," they addressed the most

serious limitations of LSM-based key–value stores for SSDs. With its optimized memory use, guaranteed maximum delays, and hardware accelerators for offloading certain sorting tasks from the CPU, PinK represents a novel and effective take on data storage for SSDs in data centers. Professor Sungjin Lee, who led the study, remarks: "Key–value store is a widely used fundamental infrastructure for various applications, including Web services, artificial intelligence applications, and cloud systems. We believe that PinK could greatly improve the user-perceived performance of such services."

So far, experimental results confirm the performance gains offered by this new implementation and highlight the potential of letting storage devices compute some operations by themselves. "We believe that our study gives a good direction of how computational storage devices should be designed and built and what technical issues we should address for efficient in-storage computing," Prof Lee concludes.

More information: Im et al., PinK: High-speed In-storage Key-value Store with Bounded Tails, *Proceedings of the 2020 USENIX Annual Technical Conference* (2020). www.usenix.org/system/files/atc20-im_0.pdf

Provided by Daegu Gyeongbuk Institute of Science and Technology (DGIST)