Storage tool adapts to what its datasets' users want to search

August 11 2020, by Adam Conner-Simons

One of the biggest challenges in computing is handling a staggering onslaught of information while still being able to efficiently store and process it. Credit: Massachusetts Institute of Technology

Big data has gotten really, really big: By 2025, all the world's data will add up to an estimated 175 trillion gigabytes. For a visual, if you stored that amount of data on DVDs, it would stack up tall enough to circle the Earth 222 times.

One of the biggest challenges in computing is handling this onslaught of information while still being able to efficiently store and process it. A team from MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) believe that the answer rests with something called "instance-optimized systems."

Traditional storage and database systems are designed to work for a wide range of applications because of how long it can take to build them—months or, often, several years. As a result, for any given workload such systems provide performance that is good, but usually not the best. Even worse, they sometimes require administrators to painstakingly tune the system by hand to provide even reasonable performance.

In contrast, the goal of instance-optimized systems is to build systems that optimize and partially re-organize themselves for the data they store and the workload they serve.

"It's like building a database system for every application from scratch, which is not economically feasible with traditional system designs," says MIT Professor Tim Kraska.

As a first step toward this vision, Kraska and colleagues developed Tsunami and Bao. Tsunami uses machine learning to automatically re-organize a dataset's storage layout based on the types of queries that its users make. Tests show that it can run queries up to 10 times faster than state-of-the-art systems. What's more, its datasets can be organized via a series of "learned indexes" that are up to 100 times smaller than the indexes used in traditional systems.

Kraska has been exploring the topic of learned indexes for several years, going back to his influential work with colleagues at Google in 2017.

Harvard University Professor Stratos Idreos, who was not involved in the Tsunami project, says that a unique advantage of learned indexes is their small size, which, in addition to space savings, brings substantial performance improvements.

"I think this line of work is a paradigm shift that's going to impact system design long-term," says Idreos. "I expect approaches based on models will be one of the core components at the heart of a new wave of adaptive systems."

Bao, meanwhile, focuses on improving the efficiency of query optimization through machine learning. A query optimizer rewrites a high-level declarative query to a query plan, which can actually be executed over the data to compute the result to the query. However, often there exists more than one query plan to answer any query; picking the wrong one can cause a query to take days to compute the answer, rather than seconds.

Traditional query optimizers take years to build, are very hard to maintain, and, most importantly, do not learn from their mistakes. Bao is the first learning-based approach to query optimization that has been fully integrated into the popular database management system PostgreSQL. Lead author Ryan Marcus, a postdoc in Kraska's group, says that Bao produces query plans that run up to 50 percent faster than those created by the PostgreSQL optimizer, meaning that it could help to significantly reduce the cost of cloud services, like Amazon's Redshift, that are based on PostgreSQL.

By fusing the two systems together, Kraska hopes to build the first instance-optimized database system that can provide the best possible performance for each individual application without any manual tuning.

The goal is to not only relieve developers from the daunting and

laborious process of tuning database systems, but to also provide performance and cost benefits that are not possible with traditional systems.

Traditionally, the systems we use to store data are limited to only a few storage options and, because of it, they cannot provide the best possible performance for a given application. What Tsunami can do is dynamically change the structure of the data storage based on the kinds of queries that it receives and create new ways to store data, which are not feasible with more traditional approaches.

Johannes Gehrke, a managing director at Microsoft Research who also heads up machine learning efforts for Microsoft Teams, says that his work opens up many interesting applications, such as doing so-called "multidimensional queries" in main-memory data warehouses. Harvard's Idreos also expects the project to spur further work on how to maintain the good performance of such systems when new data and new kinds of queries arrive.

Bao is short for "bandit optimizer," a play on words related to the so-called "multi-armed bandit" analogy where a gambler tries to maximize their winnings at multiple slot machines that have different rates of return. The multi-armed bandit problem is commonly found in any situation that has tradeoffs between exploring multiple different options, versus exploiting a single option—from risk optimization to A/B testing.

"Query optimizers have been around for years, but they often make mistakes, and usually they don't learn from them," says Kraska. "That's where we feel that our system can make key breakthroughs, as it can quickly learn for the given data and workload what query plans to use and which ones to avoid."

Kraska says that in contrast to other learning-based approaches to query

optimization, Bao learns much faster and can outperform open-source and commercial optimizers with as little as one hour of training time.In the future, his team aims to integrate Bao into cloud systems to improve resource utilization in environments where disk, RAM, and CPU time are scarce resources.

"Our hope is that a system like this will enable much faster query times, and that people will be able to answer questions they hadn't been able to answer before," says Kraska.

  More information: Ding et al., Tsunami: A Learned Multi-dimensional Index for Correlated Data and Skewed Workloads. arXiv:2006.13282v1 [cs.DB]. arxiv.org/pdf/2006.13282.pdf

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Storage tool adapts to what its datasets' users want to search (2020, August 11) retrieved 2 May 2024 from https://techxplore.com/news/2020-08-storage-tool-datasets-users.html