

Optimized software-controlled solid-state drive for big data processing

November 19 2015




Figure 1. Memory expansion for in-memory database

Fujitsu Laboratories today announced that it has developed a solid-state drive (SSD) in which flash memory can be directly controlled by software running on a server. By optimizing data positioning for access from an in-memory database, a world's first, it achieved a processing performance three times faster than that of ordinary SSDs. With in-memory databases, which enable high-speed analysis by loading data onto servers on a type of memory called DRAM (Dynamic Random Access Memory), when the volume of data exceeds the capacity of memory, lags in access to storage and other factors reduce processing speeds. As a result, there has been a desire for technology that could

expand memory using high-speed SSDs.

Now Fujitsu Laboratories has developed a SSD that enables read/write commands for each [flash memory](#) chip directly from software. It also developed a "read-ahead" feature that, in accordance with the access pattern, enables parallel data retrievals from multiple flash memory chips without the instructions interfering with each other. By loading data into DRAM before using the data from the in-memory database, data usage and data loading are processed simultaneously, enabling high-speed big data processing, with no access lags, even with limited DRAM capacity. Details on this technology will be presented at the 27th Annual Computer Systems Symposium (ComSys 2015), which is scheduled to be held on November 25 at Ochanomizu University in Tokyo.

Background

In recent years, the amount of data processed by computers, including a variety of data from sensors, is increasing. Accordingly, there is a rising need for high-speed processing technology for big data analytics, which generates new value from these massive volumes of data. With in-memory databases, which achieve high-speed analysis by storing all data in DRAM on servers, large volumes of data that exceed DRAM capacity simply cannot be processed. One way of resolving this issue is to expand memory capacity by using high-capacity SSDs, and progress is being made on technologies for separately using DRAM and memory expansion.

The performance of memory expanded with SSDs depends on the performance of the SSD itself, and because this impacts analytical performance, there has been a desire for faster SSDs.

Issues

SSDs are equipped with multiple flash memory chips, but from the perspective of software running on a server, the inside of an SSD is a black box, and read/write operations cannot be directly performed on the flash memory inside. As a result, there are times when read/write instructions are competing for the same [flash memory chip](#) which is a factor that degrades performance.

About the Newly Developed Technology




Figure 2. Newly developed software-controlled SSD

Fujitsu Laboratories has developed an SSD that enables read/write commands for each flash memory chip directly from software running on a server. It has also developed software that distributes read instructions from an in-memory database among multiple flash memory chips to operate in parallel. Features of this technology are described below.

1. Software-controlled SSD that enables read/write commands directly to multiple flash memory chips from software running on a server

In a prototype, an interface called PCI Express was used, which enables high-speed data transmissions and is installed on typical motherboards. To fully utilize the wide bandwidth of PCI Express, Fujitsu Laboratories developed a software-controlled SSD, equipped with 16 control channels and 256 flash memory chips, in which the interface enables read/write commands directly to each flash memory chip from server software (figure 2). As a result, transmission speeds of roughly 5.5 gigabytes per second were achieved.

2. Technology for parallel "read-ahead" from multiple flash memory chips prior to using data from the in-memory database




Figure 3. Accelerating the speed of an in-memory database

Fujitsu Laboratories also developed software that analyzes the memory access patterns of an in-memory database and, instead of retrieving the data from flash memory each time, retrieves the data in parallel from multiple flash [memory chips](#) in advance so that the data are read instantaneously, without competing for access.

Results

With the newly-developed software-controlled SSD and technology for data "read-ahead," Fujitsu Laboratories was able to improve processing performance by roughly three times compared to a conventional SSD, even with limited DRAM capacity. This technology therefore enables high-speed analysis of the increasing volume of [big data](#).

Fujitsu Laboratories plans to increase the parallelization of data retrievals and consider applications beyond in-memory databases with the aim of bringing this technology into practical use in fiscal 2017.

Provided by Fujitsu

Citation: Optimized software-controlled solid-state drive for big data processing (2015, November 19) retrieved 26 April 2024 from <https://phys.org/news/2015-11-optimized-software-controlled-solid-state-big.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.