

Fujitsu develops technology to automatically resolve performance problems in distributed storage

July 30 2012




Figure 1: Distributed storage

Fujitsu Laboratories today announced the development of a technology that automatically resolves problems caused by concentrated accessing of popular data items in a distributed storage system in order to curb access-time slowdowns.

Distributed storage combines multiple servers into a single storage. Increasing the number of servers improves [storage capacity](#) and performance, making this approach appropriate for storing data that

grows day by day. Furthermore, storing replicas of the same data simultaneously on multiple servers increases data reliability and access performance. If, however, there is a sharp increase in accesses to a particular stored data item, the load on the server storing it will increase, which may greatly cause an upturn in user access times.

Fujitsu Laboratories has developed a technology that can instantly detect spikes in popularity for a data item and automatically increase the number of replicas of it to level-off server loads. This automates what had previously been a manual approach for dealing with popularity spikes, and can limit slowdowns in access times. The new technology applies to distributed object storage, and in test cases of access concentrations on the Internet, has been proven to ease those concentrations by approximately 70%, resulting in an improvement in access times of tenfold or greater.

This technology makes it possible to stabilize operations on ICT systems where access patterns are difficult to predict.

Details of this technology are being presented at the Summer United Workshops on Parallel, Distributed and Cooperative Processing (SWoPP 2012 Tottori), opening August 1 (Wednesday) in Tottori, Japan.

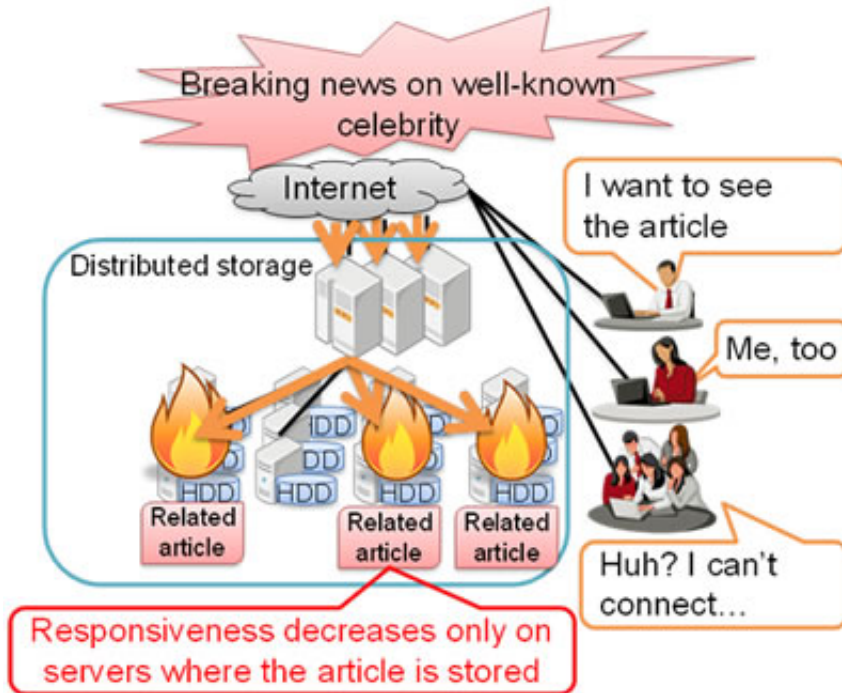


Figure 2: Performance slowdown under concentrated access of popular data item

With the popularity of smartphones and sensors, the volume of data being stored and analyzed is growing rapidly, bringing about new business value. Distributed storage is often used to store massive volumes of data, in which multiple hard drives, solid-state drives, and other storage mechanisms are combined to be treated as a single storage. Storage capacity and performance can be improved by adding more servers. In addition, simultaneously storing replicas of a single data item on multiple servers increases the reliability of data (Figure 1).

In distributed storage, concentrated accessing of a popular data item can result in a situation in which performance may have no correlation to the number of allotted servers. For example, if there is a news story with widespread public interest, most of the load will be concentrated on the server storing that popular data item, so the total number of servers is

irrelevant to performance (Figure 2).

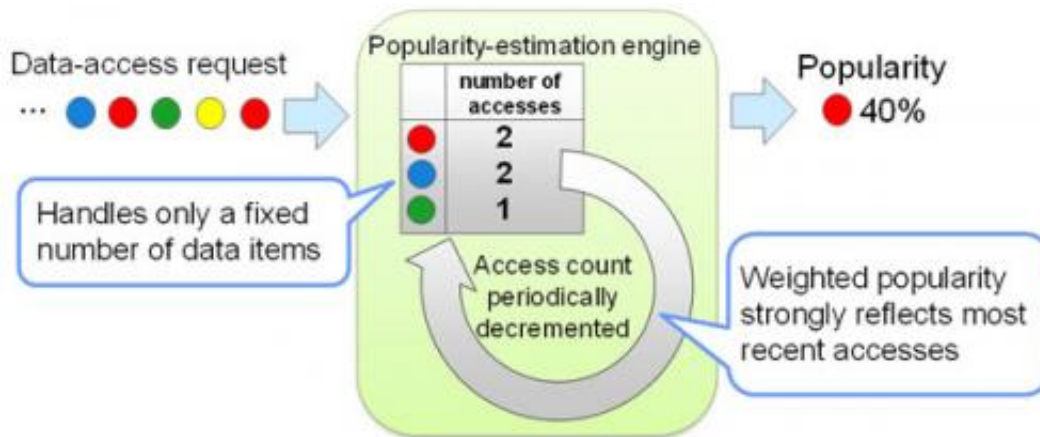


Figure 3: Popularity-estimation engine

Fujitsu Laboratories has developed a technology called Adaptive Replication Degree, which automatically detects concentrated access of a popular data item and increases the number of servers with replicas of it to distribute accesses. This technology can very rapidly detect and resolve access spikes and prevent access-time slowdowns, enabling consistently stable access performance. It even stabilizes ICT systems where access patterns are difficult to predict. Adaptive Replication Degree automatically handles the process of detecting concentrated accesses and varying the number of replicas, so the process requires no manual intervention.

Details of the technologies in Adaptive Replication Degree are as follows.

1. Using a small amount of memory to rapidly detect data items which sharply increase in popularity, causing access spikes among massive data

volumes

To detect sudden access concentrations, Fujitsu Laboratories developed a popularity-estimation engine that estimates popularity weighted by how recently it was accessed (weighted popularity) using a small amount of memory (Figure 3).

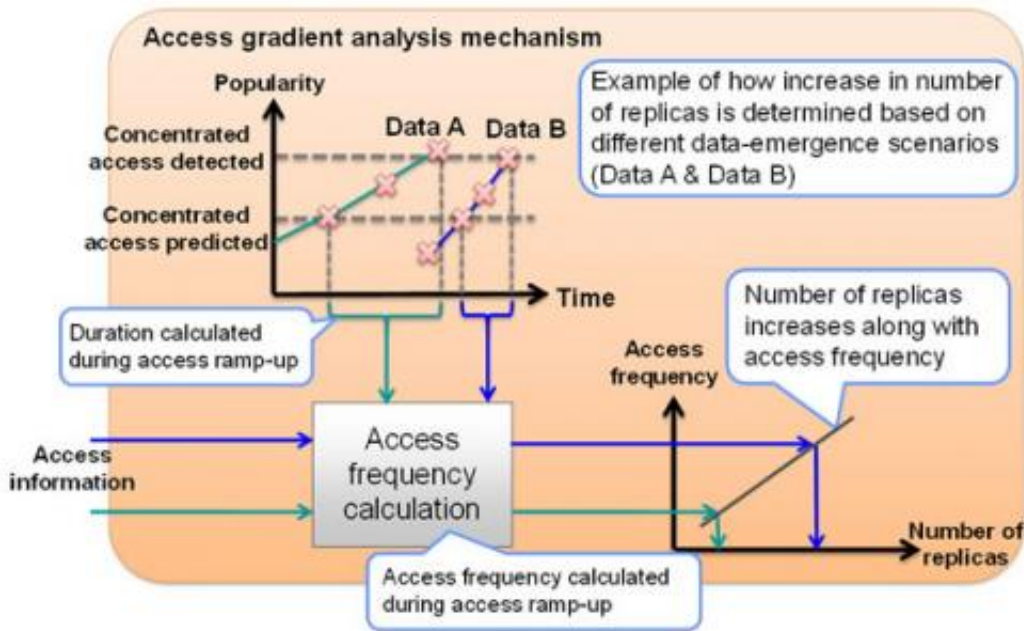


Figure 4: Access gradient analysis mechanism

The newly developed method logs the number of data accesses only for a fixed number of data items, requiring little memory. When an unlogged data item is accessed, only a minimal access count is replaced. By carrying over the access count, popularity can be estimated with a high degree of accuracy.

Also, by reducing the access counts for every fixed number of accesses,

the most recent access is weighted in the count, so that drastic changes in popularity can be detected.

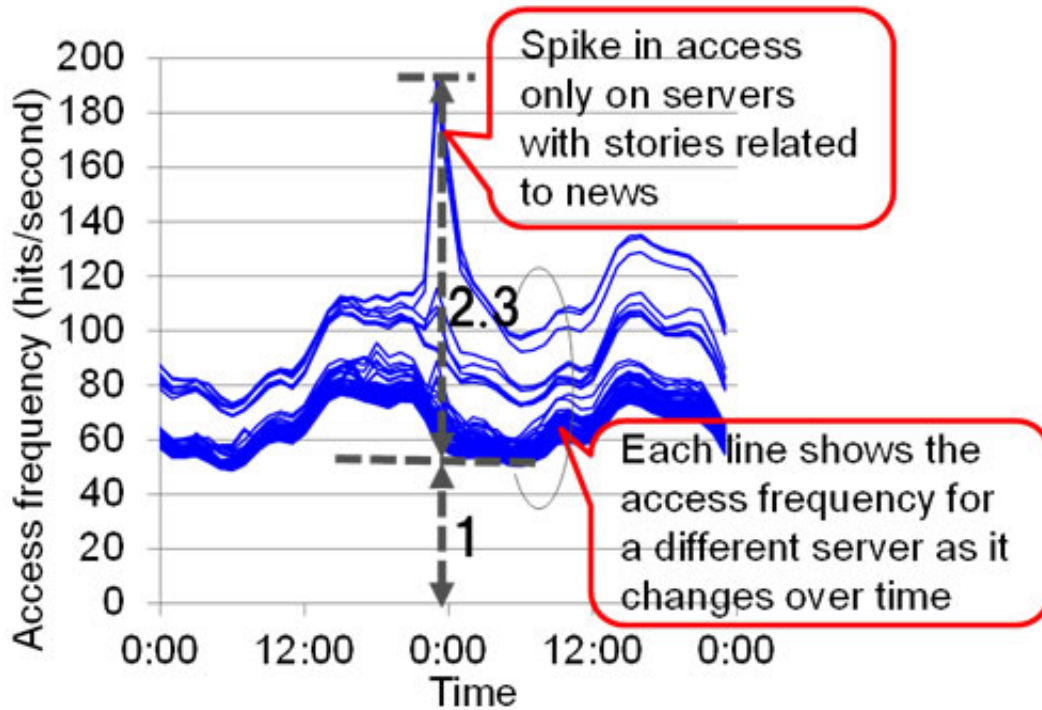


Figure 5: Changes in access frequency on each server during access concentration under previous method

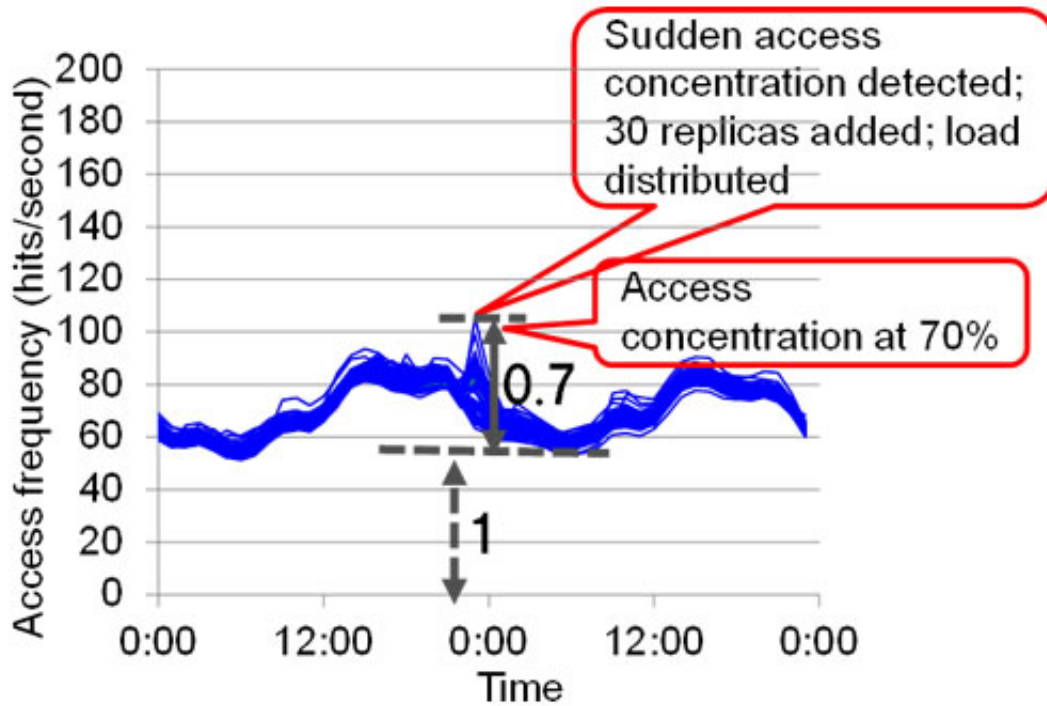


Figure 6: Changes in access frequency on each server during access concentration using this new technology

2. Automatic optimization of the number of replicas

Fujitsu Laboratories developed a technique that analyzes how highly concentrated access occurs, which causes the number of replicas to fluctuate, based on the frequency of access during periods of concentrated access (Figure 4). The number of replicas to add is automatically determined when a popular data item is detected. This technique uses two threshold values—one indicating that a spike is underway, another as a sign of an impending spike—to detect periods of concentrated access. The higher the access frequency during those periods, the more replicas are created, so that replicas are increased in proportion to the magnitude of a traffic spike.

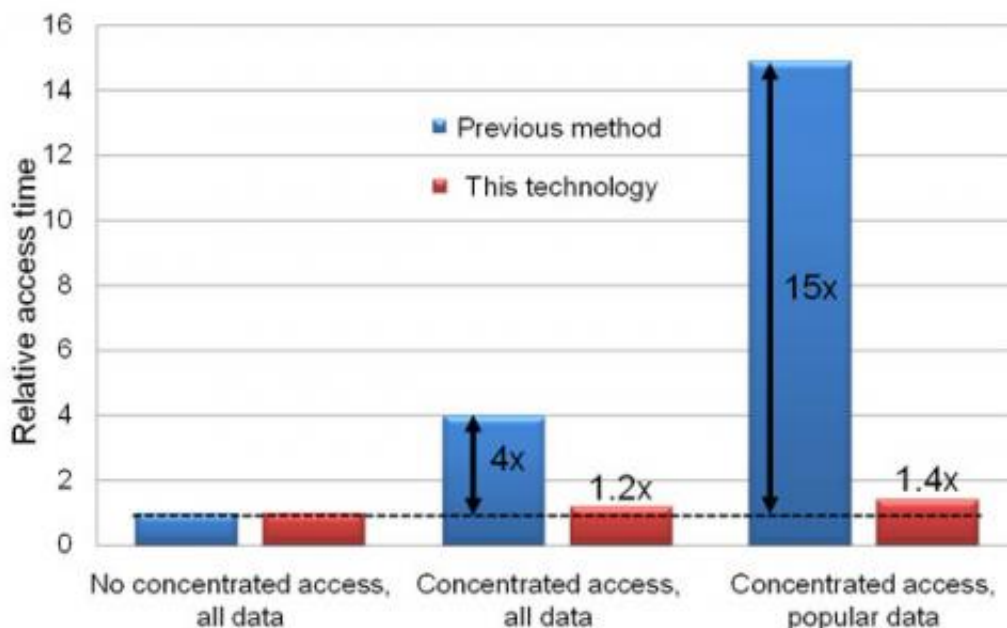


Figure 7: Relative access times for each method for different data categories, under different load conditions

This technology was tested on 64 servers in a recreation of a real-world traffic spike caused by a breaking story on the Internet about a well-known pop star (Figures 5, 6). The changes in access frequency to each server per hour were measured, and using existing methods, accesses were concentrated on only those servers that contained related data. The access frequency increased by a factor of approximately 2.3 times. Using the new technology, however, the increase in access frequency increase actually dropped to 70% of previous levels, demonstrating that this technology effectively levels off loads.

Another test was conducted using 16 servers to determine access time as viewed from the user's perspective. Figure 7 shows average access time for all data in periods of concentrated access and average access time for

popular data in periods of concentrated access. Comparing previous methods with this new technology, all access times for all data in periods of normal access were set against an average. Looking at average access time for all data, access times increased approximately by a factor of 4 times under concentrated access compared to normal access using previous methods, but only approximately by a factor of 1.2 times using the new technology. Looking at popular data, access times increased approximately by a factor of 15 times using previous methods, but only approximately by 1.4 times using this technology.

Fujitsu Laboratories is continuing to test and improve the performance of this technology with the goal of applying it to products and services in 2013.

Source: Fujitsu

Citation: Fujitsu develops technology to automatically resolve performance problems in distributed storage (2012, July 30) retrieved 6 May 2024 from <https://phys.org/news/2012-07-fujitsu-technology-automatically-problems-storage.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.