

論文の内容の要旨

Research on High Performance Database Management Systems with Solid State Disks

(SSD を用いた高性能データベースシステムに関する研究)

氏名 王 永 坤

In this information explosion era, data volumes grow drastically, posing great challenge to the data-intensive applications, such as the database management systems. These data-intensive applications are required to process the huge amount of data quickly. However, in the current hard disk-based storage system, the speed gap between the CPU and hard disks becomes the bottleneck to improve the performance. At this time, the Solid State Disk (SSD) is on the spotlight. The SSD, mainly composed of flash memory, has a significant performance advantage over the traditional hard disk. The read performance of SSD is about two orders of magnitude better than that of hard disk. The sequential write performance of SSD is also much better than that of hard disk. However, the random write performance of SSD is comparable or even worse than that of hard disk, because of the "erase-before-write" design of the flash SSD. Therefore, comprehensive study is required to incorporate the flash SSDs into the existing enterprise database management systems.

In this thesis, I performed a research on the possibility of building high performance database management systems with SSDs. Firstly I provided the basic performance study of the flash SSD. I built a micro benchmark to bypass the operating system buffer cache to get the real performance of flash SSD. Based on the micro benchmark, I also designed a modeling and simulation tool, which is expected to be effective to design a large storage system with flash SSDs. Secondly I studied the performance benefit by incorporating a special technique, the Non-In-Place-Update technique into the IO path. With this technique, the performance of database management systems built on the flash SSDs is improved significantly. Next, I studied the IO behavior along the IO path of flash-based database system, with different IO strategies, different database applications, different enterprise-class SSDs. I then proposed to use the application hint to improve the performance further. I presented the experimental evaluation and analysis. Finally, I conclude the usage of the flash SSD in the database management systems.