Online Monitoring of Database Structural Deterioration

Takashi Hoshino
Graduate School of Information Science and Technology,
University of Tokyo
hoshino@tkl.iis.u-tokyo.ac.jp

Kazuo Goda and Masaru Kitsuregawa
Institute of Industrial Science,
University of Tokyo
{kgoda,kitsure}@tkl.iis.u-tokyo.ac.jp

Abstract

The paper proposes a structural deterioration monitor of database, which is an essential building block to realize an autonomic database reorganizer. Experimental results with our prototype show that the monitor can keep track of structural deterioration with high resolution, high accuracy, and low overhead.

1 Introduction

The physical/logical data structure of database may gradually deteriorate as updates are repeatedly performed on the database. This phenomenon is called *structural deterioration*[1], which can degrade data access performance. *Database reorganization* [2, 3] removes the structural deterioration by relocating data in secondary storage to recover performance. Accordingly, database reorganization is an essential task for especially the systems which require high performance and high availability.

Database reorganization generally consumes much time to complete, issuing a massive number of IOs, thus it is likely to impact on online workloads. Database administrators must carefully schedule a database reorganization task in order to manage system performance; specifically, the administrators should determine the portion of the database and the points of time to be reorganized, by considering structural deterioration and reorganization cost. Precise up-to-date quantitative information of the database structural deterioration is crucial to derive an efficient reorganization plan, however, so far such a monitoring facility was not investigated by the research community or industry. To date, the database administrator tends to make reorganization plans based on rules of thumb by using mainly measured performance and rough statistics of the database. This kind of naive solutions leads to inefficient and expensive database administration.

The paper proposes an online monitor of database structural deterioration as a first research step for autonomic database reorganization. The monitor has the capability of keeping track of the up-to-date structural deterioration with high resolution, high accuracy, and low overhead. We developed a GUI tool for visualizing the measured structural deterioration in a real-time fashion, so that the database administrator can easily understand the deteriorated portion of the database and the degree of performance impact.

2 Architecture

We describe a grand design of our proposed online monitor of database structural deterioration. The monitor consists of a *sniffer* and an *estimator*. The sniffer implemented in low-level storage engine of database systems, captures events of structural changings and sends the events to the estimator. Suppose B+tree structure, which is often used for typical database systems. In a B+tree structure, given that many records are inserted into a particular page and the page becomes full, the page splits; specifically, a new page is allocated and several records of the old page move to the new page. Such a 'page split' event is captured by the sniffer and the structure-related information such as page id, associated key range is sent to the estimator. Upon the received information, the estimator calculates the degree of the structural deterioration for a considered access on the basis of structure model[1]. Here the degree of structural deterioration is defined as a deteriorated ratio of access performance. Let us suppose a typical range scan of B+tree. If the scan takes 300 seconds in the well-organized database and 450 seconds in the deteriorated database, then the deterioration degree becomes 1.5.

The monitoring facility has three technical advantages. First, the 'structural changing' event can be captured for each concerned page and and the structural deterioration can be estimated in the unit of page. That is, fine-grained monitoring of the structural deterioration is realized, facilitating administrators' identification of the deteriorated portion of the database. Second, physical performance char-

[1] We have to omit the discussion of the model due to the page limitation.

Fourth International Conference on Autonomic Computing (ICAC'07)
0-7695-2779-5/07 $20.00 © 2007 IEEE

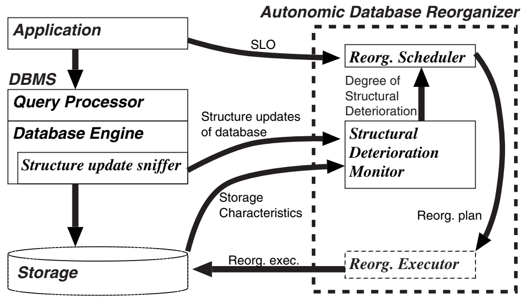

Figure 1. Design Overview of Autonomic Database Reorganizer

acteristics of the storage device (e.g. hard disk drive) are considered in calculating the deterioration degree. Thus, highly accurate estimation of the structural deterioration is expected. Finally, the sniffer need to capture a small number of events and the estimator can update the measured structural deterioration in a incremental fashion, instead of full table scanning. The additional cost of online deterioration monitoring can be very small.

We developed a framework of autonomic database reorganization by the use of the above proposed structural deterioration monitor. In the framework, as illustrated in Figure 1, a reorganization scheduler can build an efficient reorganization plan based on the measured structural deterioration degree and application specific performance requirements, and then a reorganization executor can reorganize the database based on the plan.

3 Evaluation

We implemented a prototype of the proposed monitor for MySQL (InnoDB storage engine) on Linux operating system. Figure 2 shows an example of structural deterioration visualized by the developed GUI tools of the prototype.

Experiments were conducted to evaluate the prototype. First, we analyzed the accuracy of the monitor by comparing the execution time and the measured deterioration degree of gradually refreshed LINEITEM table of TPC-H benchmark. The measured error was at most 6%. Second, we checked the resolution of the monitor under practically enough accuracy by doing the same experiment for different sizes of scan range. The measured error was at most 15% in most cases. Finally, we verified the performance overhead of the sniffer by using two types of synthetic workloads. The measured performance overhead was 0.5% for 'data load' workload and 3.2% for 'mixed bulk deletes and inserts' workload. Therefore, the above experiments show that the deterioration can be monitored with high accuracy,

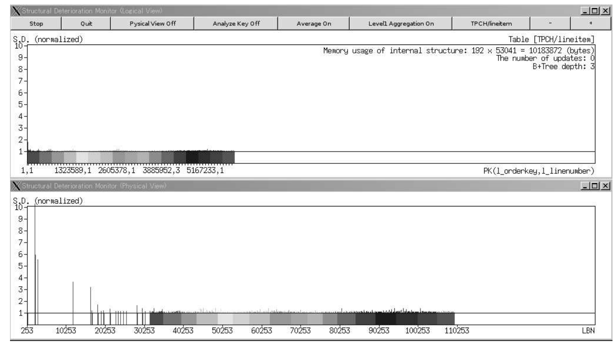

(a) after data load

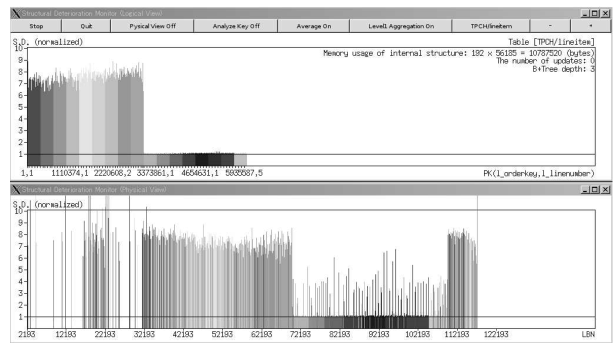

(b) after 50 refreshes

Figure 2. Structural deterioration view of LINEITEM table with S.D. monitor

high resolution and low overhead.

4 Conclusion

The paper proposes a real-time online structural deterioration monitor, an essential function for autonomic database reorganization. The monitor can measure structural deterioration of database online with high accuracy, high resolution, and low overhead. The monitor facilitates structural deterioration analysis and enables efficient reorganization planning, accordingly relieving the burden of database administrators. We would like to develop a reorganization scheduling mechanism in future work.

References

[1] T. Hoshino, K. Goda, and M. Kitsuregawa. A Study on Performance Analysis of Structural Deterioration for RDBMS. In *SMDB Workshop*, pages 81–90, 2005.

[2] G. H. Sockut and R. P. Goldberg. Database Reorganization - Principles and Practice. *ACM Comput. Surv.*, 11(4):371–395, 1979.

[3] (Ed.)D.Lomet. Special Issue on Online Reorganization. *IEEE Data Eng. Bull.*, 19(2):1, 1996.

Fourth International Conference on Autonomic Computing (ICAC'07)
0-7695-2779-5/07 $20.00 © 2007 IEEE