

Scalability and Resource Usage of an OLAP Benchmark on Clusters of PCs

Michela Taufer, Thomas Stricker, Roger Weber

Department of Computer Science
ETH Zentrum
CH-8092 Zurich, Switzerland
taufer,tomstr@acm.org
weber@inf.ethz.ch

Abstract

Designing clusters of PCs for distributed databases processing OLAP (On Line Analytical Processing) workloads in parallel with good scalability remains a particular challenge as we are lacking a deep understanding of the architectural issues around resource usage by standard DBMSs on distributed platforms.

To address this problem, we present a novel performance monitoring framework for filtering and abstracting samples of performance data from low level counters into a high level performance picture. Our framework is used side by side with the DBMS and delivers many interesting insights about the most critical resource in the different queries and systems configuration. As required for a larger distributed hardware/software system, our solution comprises software instrumentation at the OS level, tools for gathering performance relevant data and an analytical model for performance evaluation and performance prediction to future platforms.

We demonstrate the viability of our approach with the in-depth analysis of distributed TPC-D, a standard OLAP benchmark running on clusters of commodity PCs. Based on the data

provided by our framework, we isolate and resolve a few crucial performance issues of OLAP workloads on clusters. For different queries, we give a workload characterization in terms of resource usage, quantify the optimal scalability and investigate the impact of the networking speed on the overall application performance. We show that the disk performance and CPU speed remains the most critical resource bottleneck for most queries. Queries with a lot of inter-node communication are limited by the communication software inefficiency within the DBMS and not by the raw networking speeds. A systematic performance evaluation constitutes a solid basis for architectural decisions and system optimization in clusters of PCs that are dedicated to large parallel database systems.

Keywords: *parallel databases, distributed OLAP processing, cluster of PCs, performance analysis, workload characterization.*