Energy Cost, The Key Challenge of Today's Data Centers:

A Power Consumption Analysis of TPC-C Result

Introduction

The paper elaborates about making a power consumption model benchmark based on TPC-C full disclosure reports on a various transactional systems. It analyses how the trends are going in terms of performance and power consumptions on TPC-C workloads with a defined system configurations. Ultimately, the paper is trying to reveal the potential solutions that could be applied according to the result of the experiments.

The Workload Constraints

The TPC-C workload basically resembles an online transaction processing (OLTP) workload. It is a mixture of read-only and update-intensive transaction which evaluates key performance factors such as user interface, communications, disk I/Os, data storage, backup, and recovery. The goal of this workload is to maintain the system components to be equally fully utilized to get an upper bound of how big the system can exerts performance against power consumptions. However, the workloads of the TPC-C are not used fully. The experiment reduces the amount of queries to a certain level although it is being said that it still retains its essential performance characteristics. We are forced to believe that there are no errors rate or deviation rate that caused by this decision since they provide no proof on how they reduce the workloads.

The measurement of the TPC-C workloads is using a price-performance metric: the cost of three-year ownership of all components divided by transaction per minute (tpmC). The cost of ownership of a particular component is obtained via an official website of benchmark oriented service which then raises the question of its relevancy or reliability.

Power Consumption Estimation

The paper estimates the power consumption only on Tier 2 (Client) architecture and Tier 3 (Server) architecture which is an obvious thing to do. However, in their future works, they insist that they need to include Tier 1 (Driver) architecture to be estimated its power consumption. It is rather strange that they want to include this since the driver architecture usually owned by the user not the vendor.

Since all components are being fully utilized, the peak power consumption of all system components can be derived from the sum of the peak power consumption of individual component used in the corresponding system. The peak power consumption of individual system component is obtained from the manufacturer. On the verification of the power model, it overestimates the power consumption by about 10-25%. This is a rather huge amount of overestimation which leads to the question of how good the power model really is.

Result Analysis

It is questionable that the paper uses a linear line to derive the result from the TPC-C workload while an exponential line is fit more accurately to derive the graph of transaction performance, price performance, and power consumption. The transaction performance graph is also a little bit misleading since the linear line goes below 0 tpmC which is an impossible situation to have. The result of the TPC-C also only considers a single system with a defined amount of processor which raises a question about how it will represent the result for bigger systems. Although, the paper argues this can be done by summing up all the system that is used to define the trends, it is questionable how this configuration will impact the outcome of the graph trends.

Solution Offered

The result of the TPC-C stated that storage subsystem is the main power consumer from the system being tested. The paper suggested a lot of hardware improvement on storage subsystem to mitigate the huge power consumption on storage subsystem. However, the paper fails to provide the performance impact ration that could be gain for using this decision.

Observations

The TPC-C benchmark is aimed at capturing the throughput performance of systems under test. For this reason all measurements are done while the system at peak workload (during the so-called "stationary" phase) – while the results of these measurements capture the performance of the system throughput-wise, the behavior of the system at full utilization may be irrelevant for the overall power efficiency. For instance a system may be optimized for energy consumption for a medium workload; applying the above mentioned metric to the results of the TPC-C runs would fail completely in capturing the optimizations and therefore rate the system poorly energy-wise (despite the fact that it would behave very well under normal working conditions). Arguably, the TPC-C benchmark may not provide relevant test cases when trying to rate the system in what power consumption is concerned; a better approach would be to design test cases try to mimic the normal operation of the system or at least measure the power performance at several different intermediary workloads.