Chameleon: A Self-Adaptive Energy-Efficient Performance-Aware RAID

Qingbo Zhu and Yuanyuan Zhou, Department of Computer Science, UIUC

1 Motivation

Recently, researchers have considered to reduce the energy consumption in large disk arrays. The problem is significant and compounded by the proliferation of large data centers that are heavily reliant on disk arrays to meet their storage needs. The cost of powering the nation's data centers is expected to reach $4B by 2005 with a projected increase of 25% annually. Considering that disk storage accounts for 27% of the total energy cost in data centers, designing energy-efficient disk arrays can significantly reduce the Total Cost of Ownership.

A solution that reduces energy costs in disk arrays for data centers cannot violate performance expectations that the centers provide to their customers. For example, a typical Service Level Agreement specifies the percentage of transaction response times which cannot exceed a value (e.g. seconds). The penalties for failing to comply are costly. Therefore, employing energy conservation techniques in data centers must be carefully balanced with established performance guarantees.

There is, however, a natural conflict between these two goals. A performance oriented organization suggested by a traditional RAID stripes data across multiple disks to provide uniform load balancing and thus achieves high performance. However, no single disk is idle enough to spin down. Conversely, Popular Data Concentration(PDC), a recently proposed energy efficient data layout, concentrates the load by periodically migrating data based on their popularity to allow certain disks to spin down. While potentially suitable in light workloads, it suffers serious performance penalties during heavy workloads due to the hot spots.

Moreover, data center workloads change dynamically, requiring versatile solutions that are self-adaptive to coincide with workload characteristics. However, none of current solutions can adapt to workload changes so as to dynamically achieve a balance between energy conservation and performance guarantees.

2 The Design of Chameleon

Our goal is to dynamically maximize energy conservation while still meeting performance guarantees. Specifically, we should be able to conserve as much energy as PDC during light loads and provide high performance comparable to the traditional RAID during heavy loads.

The solution we propose here is Chameleon, a self-adaptive energy-efficient performance-aware RAID that combines the performance benefits of a traditional RAID with an energy-efficient disk layout. The main features of our Chameleon are listed below:

Two-tiered Data Layout Chameleon relies on a two-tiered data layout which is automatically and transparently adapted based on an analytical performance/energy model given the observed workloads. It stores hot data in an active disk group and cold data in an inactive disk group, both using RAID-5 data organization. The disks in active group remain spinning at full speed to provide a high degree of performance while the inactive disks are set to be spinning at low speed to conserve energy (Sony reported such a prototype 2-speed disk). Chameleon dynamically adjusts the size of the two tiers and migrates data blocks between them to meet changing workload requirements.

An Analytical Performance/Energy Model Chameleon applies an analytical model to workload characteristics to determine the optimal disk configuration. This is achieved by solving an optimization problem that maximizes the number of disks that spin at low speed given performance requirements.

Fast Disk Shuffling Chameleon employs a novel technique called "randomized shuffling" to move disks quickly and efficiently between two tiers and achieve load balance within both groups. All of the operations are performed sequentially to avoid the cost of seek and rotational delays during disk reorganization.

3 Experiments

We conduct experiments by running an OLTP trace on the widely used Disksim simulator. We enhanced it with a two-speed disk model which can service requests at speeds of both 15000RPM and 3000RPM. The disk parameters are similar to those of IBM Ultrastar 36Z15.

Our preliminary results show Chameleon can conserve as much energy as PDC during light workloads and provide high performance comparable to the traditional RAID during heavy workloads. The model is pretty accurate since the difference between the estimated average response time and the measured one is not more than 5.2%. Moreover, the cost of disk shuffling is roughly as small as sequentially scanning a disk.

4 Conclusions and Future Work

Preliminary results show Chameleon is able to dynamically achieve a balance between performance expectations and energy conservation. For future work, we will first investigate the effects of multiple speeds; second, we are about to examine additional traces other than OLTP; third, we will implement Chameleon in the V3 storage system to evaluate its effect in real systems.