Storage Device Performance Prediction with CART Models

[Extended Abstract]

Mengzhi Wang, Kinman Au, Anastassia Ailamaki,
Anthony Brockwell, Christos Faloutsos, and Gregory R. Ganger

Carnegie Mellon University

ABSTRACT

This work explores the application of a machine learning tool, CART modeling, to storage devices. We have developed approaches to predict a device's performance as a function of input workloads, requiring no knowledge of the device internals. Two uses of CART models are considered: one that predicts per-request response times (and then derives aggregate values) and one that predicts aggregate values directly from workload characteristics. After training on the device in question, both provide reasonably-accurate black box models across a range of test traces from real environments. An expanded version of this paper is available as a technical report [1].

Categories and Subject Descriptors

C.4 [**Performance of systems**]: Modeling techniques; D.4.8 [**Performance**]: Modeling and Prediction

General Terms

Experimentation,Performance,Management

Keywords

Performance prediction,Storage device modeling

1. INTRODUCTION

The costs and complexity of storage management make automation of management tasks a critical research challenge. One key aspect of this, particularly for large storage infrastructures, is deciding which data sets to store on which devices. Doing so requires the ability to predict how well each device will service each workload, so that load can be balanced and good matches can be exploited.

Performance models for such prediction have long been utilized by researchers to compare alternate designs. Given sufficient effort and expertise, an accurate simulation or analytic model can be generated to explore design questions for a particular device. Unfortunately, in practice, such time and expertise is not available for deployed infrastructures. Such infrastructures often consist of numerous and distinct device types, and their administrators have neither the time nor the expertise needed to configure device models.

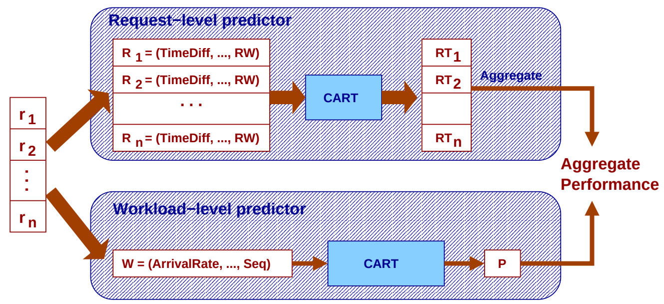

Figure 1: Per-request and feature-based predictors. r_i is the i-th request in the workload; RT_i is the response time of r_i; R_i is the set of per-request characteristics for r_i; W is the set of workload-level characteristics.

Automated, black box model generation approaches are desirable for such environments. By "black box," we mean that the model (and model generation system) has no information about the internal components or algorithms of the storage device. Given access to the device for some "training period," a model can be built to determine the performance as a function of input workload. Such black box modeling is difficult, because storage device performance is a complex and non-linear function of the input workload. In addition, storage workload characterization remains an open problem.

We employ a standard machine learning tool, Classification And Regression Tree (CART) modeling, to address the performance prediction problem. CART models provide nonlinear piecewise-constant mappings from input space to output space. Although other families of models, including neural networks, could be used, CART models are convenient for several reasons. They are easy to fit, can be represented visually by trees (leading to good interpretability), and can provide good approximations to highly nonlinear mappings.

2. PREDICTING PERFORMANCE

The goal of this work is to build a model for a given device

SIGMETRICS/Performance'04, June 12–16, 2004, New York, NY, USA.
ACM 1-58113-664-1/04/0006.

to predict the performance of a workload on the device. The model takes a sequence of disk requests as input and predicts the performance (e.g., average or the 90th percentile) of the workload. Building the models requires training on observations of the device under some workload(s). Figure 1 illustrates two approaches using CART models for device modeling. The difference between the two approaches lies in the level of detail of the input to CART models.

The first approach, a `request-level` predictor, predicts the response time of each disk request based on characteristics of the request and those before it in the sequence. As a result, the model can produce any aggregate performance metrics from the predicted response times. The second approach, a `workload-level` predictor, predicts an aggregate performance metric directly from workload-level characteristics, such as metrics of burstiness, locality, and read/write ratio.

There is a clear tradeoff between the two predictors. The `request-level` predictor is more computationally demanding than the `workload-level` predictor. The former needs to run the CART model for each request in the workload, while the latter requires only one run per workload. On the other hand, the `request-level` predictor is more robust because aggregating multiple predictions reduces the total error. In addition, the workload-level characteristics are much harder to discover and quantify.

3. EXPERIMENTAL RESULTS

Figure 2 shows the prediction accuracy of our predictors on three traces. Each trace contains four weeks of activities from a different workload. The device being modeled is a 9GB Atlas 10K disk. The average response times are 83.78, 115.71, and 259.88 milliseconds for *cello92*, *cello99-1*, and *cello99-2* respectively.

All predictors are trained on the first two weeks of *cello99-1*. The prediction accuray of *cello99-1* is measured on the last two weeks. Instead of predicting the entire workloads in one shot in the `workload-level` predictor, we chop the trace into one-minute fragments and make predictions for each fragment. This allows us to get enough training data points from two weeks of training trace.

The graphs compare the CART-based predictors with three other predictors. The `constant` predictor uses the average response time of the training trace as its prediction. The `periodic` predictor predicts all fragments of the same offset within a week to have the same performance. Note that these two predictors are not real device predictors and the input workload is not an input of them. The `linear` predictor uses a linear regression of the workload-level characteristics to produce predictions. The difference in prediction accuracy of `linear` to the CART-based `workload-level` predictor indicates the power of nonlinear models, such as CART models.

Overall, the two CART-based predictors are robust across all the workloads. Not surprisingly, `constant` and `periodic` incur sky-high prediction errors when the workload is not the same as that on which they are trained. This illustrates that they model the workloads rather than the storage device. In contrast, consistent prediction accuracy proves that the CART-based predictors are less tied to the workloads. The `request-level` predictor does extremely well, achieving a mean absolute error of about 10 milliseconds for all the workloads. The absolute error for the `workload-level` predictor, on the other hand, grows with the average and variance of the actual response time of the workload being predicted. This is due to the increased complex access patterns in workloads with a high average and variance response time.

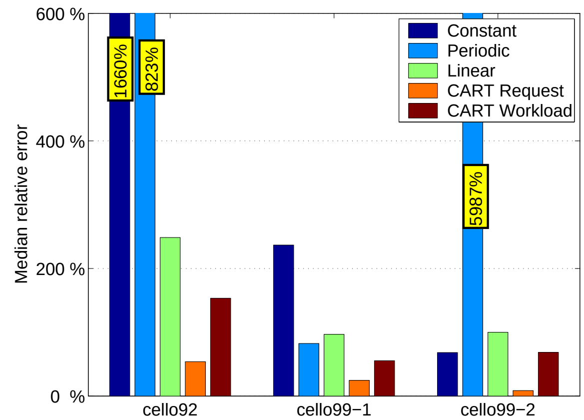

(a) Median relative error

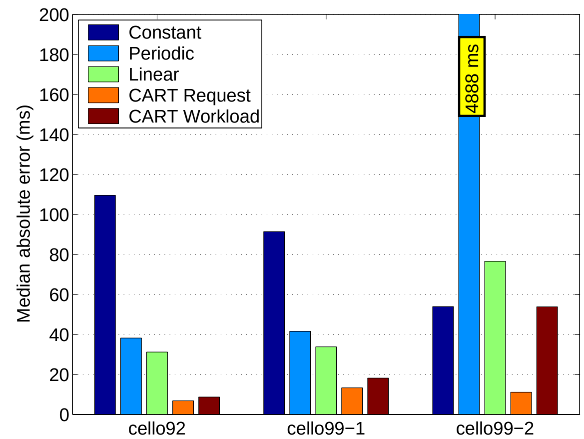

(b) Median absolute error

Figure 2: Average response time prediction accuracy. The device being modeled is a 9GB Atlas 10K disk. All predictors are trained on two weeks of *cello99-1* and tested on the workloads indicated.

4. ACKNOWLEDGMENT

We thank the members and companies of the PDL Consortium (including EMC, Hewlett-Packard, Hitachi, IBM, Intel, LSI Logic, Microsoft, Network Appliance, Oracle, Panasas, Seagate, Sun, and Veritas) for their interest, insights, feedback, and support. We thank IBM for partly funding this work through a CAS student fellowship and a faculty partnership award. This work is funded in part by NSF grants CCR-0205544, IIS-0133686, BES-0329549, IIS-0083148, IIS-0113089, IIS-0209107, and IIS-0205224.

5. REFERENCES

[1] Mengzhi Wang, Kinman Au, Anastassia Ailamaki, Anthony Brockwell, Christos Faloutsos, and Gregory R. Ganger. Storage device performance prediction with CART models. Technical Report CMU-PDL-04-103, Carnegie Mellon University, 2004.