CCIT Report #881 March 2015

Analysis of Write Amplification with Time Locality and/or Multi-Write Access

Saher Odeh
Technion EE Department
sahero@tx.technion.ac.il

Yuval Cassuto
Technion EE Department
ycassuto@ee.technion.ac.il

ABSTRACT

Write amplification is a central performance figure of solid-state storage devices. As a result, significant effort has been devoted in the storage systems community to analyzing write amplification in useful scenarios. Analysis here means that a parametrized model is chosen for the device and its incident workload, such that the expected average write amplification can be obtained mathematically. While reaching this goal necessitates a detailed mathematical argumentation, such analysis offers the great benefit of predicting performance without need to deploy or simulate the device. In this paper we develop an analysis framework to calculate the write amplification for two novel and practically important scenarios: workloads with time locality, and devices with multi-write capabilities. The former captures a central feature of real-world workloads, while the latter addresses a promising feature likely to be added to next-generation solid-state storage devices.

1. INTRODUCTION

Non-volatile storage devices have become central components in almost every information system. It is likely that the adoption of such devices will continue to grow fast, thanks to the superior access performance they offer and the impressive density scaling they sustain. In particular, NAND-flash based storage devices are currently the champions of non-volatile storage in terms of performance and density for unit cost. All their advantages notwithstanding, there are two major issues that complicate the adoption of NAND-flash storage devices in real systems. One is their non-deterministic access performance due to address indirection done at the flash translation layer (FTL). The other issue is their limited endurance, which degrades their reliability as they age in the system. Both of these issues are significant, and thus must be addressed by the device vendors during design, and monitored by the customers during operation. Nevertheless, a major challenge is that both performance variability and endurance issues are driven by complex hardware and software mechanisms in the device. Consequently, the storage community has sought deterministic and clear characterizations of these issues, which will allow evaluating and comparing different storage devices. The principal such characterization adopted by the storage industry is the *write amplification* (WA), which is defined as the average number of physical writes performed per user write. The write amplification emerged as a central performance measure, because it captures both the access-performance and endurance issues: writing more internally in general implies worse read/write rates for the user, as well as higher cell wear.

While not a perfect characterization of non-volatile storage performance, the write amplification has fostered great innovation as a tool to design better storage devices and to understand their behaviors. The most immediate use of the write amplification is to benchmark existing or newly proposed device architectures against relevant workload traces. In this usage, we run a workload trace through the architecture in discrete-event simulation, count the total number of physical writes performed by the device, and normalize by the number of user writes in the trace to get the empirical write amplification. To get a more comprehensive quantitative evaluation of the device, we complement the workload traces with synthetic workloads generated by some probabilistic model, most commonly the *uniform workload*. The advantage of synthetic workloads is that they are easier to obtain than recorded traces, and thus can test devices for longer, more realistic durations.

To make the device evaluation even more efficient, we are interested in methods to derive the write amplification analytically, obviating the need to run long and resource-consuming simulations. Indeed, several works in the literature successfully provide analytical tools to estimate the write amplification exhibited by a storage device given its design, e.g. the garbage collection policy it employs, and the workload that it serves. In the sequel we survey some of these works. In the mean time we note that in general the write-amplification analysis task becomes very difficult very fast as we attempt to capture more realistic device architectures and workload models.

1.1 Our Contributions

The study reported in this paper contributes several novel write-amplification analysis tools for important workload models and device architectures. Specifically, we list the following items as our main contributions

1. Defining a new model for workloads with *time locality*.