September 20, 2017

SSD Admission Control for Content Delivery Networks

Dragos Ionescu

Richard Schooler

Kenneth Barr

SSD Admission Control for Content Delivery Networks

ABSTRACT

Content delivery networks (CDNs) are utilized to deliver content with low latency. Solid State Drive (SSD) devices are a cost-effective option for intermediate tier caches in servers used in CDNs. However, excessive writes of content to SSD devices can degrade useful life of such devices. This disclosure describes admission control techniques to manage content in an SSD cache. In particular, a popularity metric for content items is determined based on an inter-arrival time. The popularity metric takes into account past admission history and is dynamically to account for variations in content popularity over different time periods. These techniques balance the SSD cache write rates to optimize the device life and cache performance.

KEYWORDS

- admission control
- popularity metric
- inter-arrival time
- SSD cache
- CDN

BACKGROUND

A Content Delivery Network (CDN) provides online content such as video to a large number of users distributed across geographic locations. An important objective for a CDN is to provide low latency such that the duration of time between receiving a content request and serving the content is minimized. Due to limitations imposed by the speed of light and other transmission delays, CDN servers are located across multiple geographies, e.g., within a few

hundred miles of user locations, to achieve latency of the order of tens of milliseconds. Use of multiple servers also provides higher content delivery capacity.

Geographically distributed servers that are part of a CDN are often known as “caches” and are configured to provide copies of the content. A cache server typically includes a multi-tiered cache, with a very fast first-level cache, e.g., that uses relatively expensive DRAM (Dynamic Random-Access Memory), and a larger second-level cache that uses much cheaper but slower hard disk drives (HDD). Some cache servers also utilize an intermediate cache tier of SSD (Solid-State Drive).

SSDs lie between DRAM and HDD in terms of cost, performance, and capacity. SSDs have high peak read and write throughput rates and relatively low “write endurance.” A typical SSD can sustain only about five device writes per day (DWPD) if it is to last a few years. This is often not a problem, as SSDs are often used in read-mostly environments. In a CDN environment, an excessive write rate to the SSD tier can lead to early wear of the hardware, resulting in reduced usable life and consequently, higher cost.

The content stored in servers that are part of a CDN typically follows a popularity distribution, such that there is a relatively small amount of popular content that serves the majority of requests (“active content”) and a long tail of less popular content that is infrequently requested. If the active content fits into the SSD capacity, the write endurance problem is mitigated. However, if the active content set is large, e.g., in the case of video data, the active content per server can be of the order of hundreds of terabytes (TB), too large to fit in a typical SSD cache tier. In such conditions, the SSD cache is subject to frequent churn, as newly-popular content is admitted and written to the SSD cache. Frequent rewrites lead to reduced usable life for the SSD.

DESCRIPTION

This disclosure describes admission control techniques that control the content admitted or written to a SSD cache. Implementing the techniques ensures that write endurance is not exceeded, while maintaining the benefit of the SSD read performance.

A popularity database stores the popularity for each distinct piece of content, e.g., a specific video, or a specific video rendered in a specific format. The database records a frequency of access for each piece of content. In a simple scheme, the time of previous access of the content is stored, e.g., using a last-access timestamp. This provides an inter-arrival time metric. This scheme can be extended to store more detailed past access patterns. To ensure an accurate measure of popularity, repeated accesses by the same user or device are discounted. For instance, in a streaming video application, subsequent pieces of the same video are accessed by the same client. Such accesses are not counted as contributions to the popularity of the video. The database also stores other details regarding the last access.

Fig. 1 illustrates an admission control technique. The popularity metric for the content, e.g., the inter-arrival time, is computed over an appropriate demand cycle (102). If the content has not been observed before, the content is assumed not popular and is not admitted to the SSD cache. If the content has been observed before, the elapsed time between the observations is the inter-arrival time (IAT).

If the inter-arrival time is less than a threshold such that the popularity metric exceeds the threshold (104), the content is determined as popular and admitted to the SSD cache (106). If the popularity metric does not exceed the threshold, the content is not admitted to the SSD cache (108). The admission history is updated (110). To prevent the size of the database from

growing without bound, the database is size-limited by discarding the oldest elements and by occasional purging.

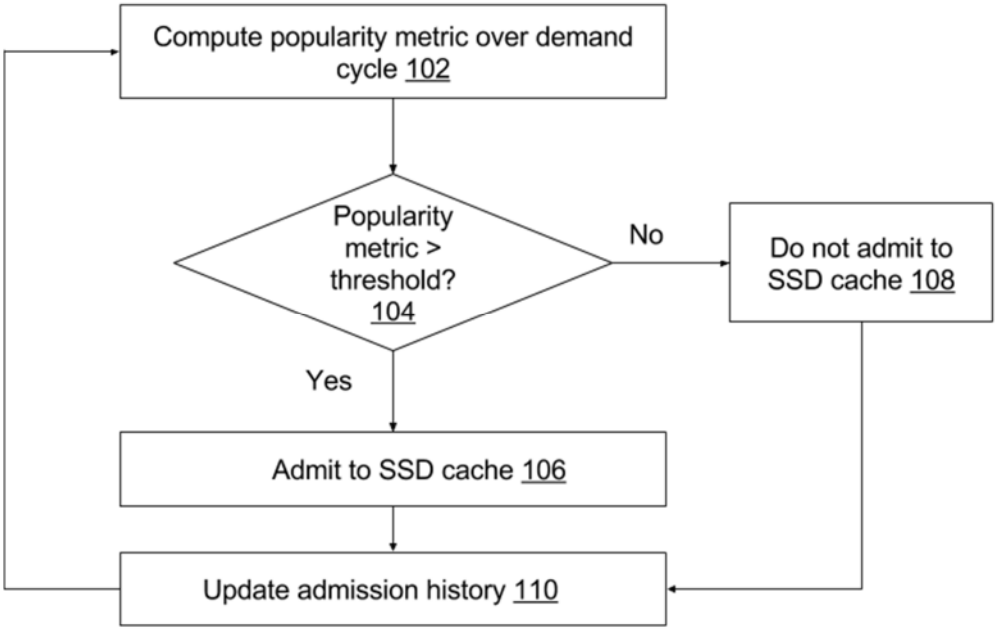

Fig. 1: Admission control

Determination of an appropriate inter-arrival threshold is important for admission control. A simple scheme is to not admit content that is less popular than the least popular content already in the cache. However, such a policy can still lead to admission of too much content and exceeding the SSD write endurance requirement.

A more aggressive policy that dynamically adapts the admission threshold is used. Per this policy, if too much content is being admitted, the threshold is tightened. For example, if a smaller inter-arrival time is observed in recent content requests, leading to a greater than permitted rate of admission of content to the SSD cache, the inter-arrival time threshold is lowered, lowering the amount of content that is admitted. Similarly, if the amount of content being admitted is below that permitted by the write endurance constraint, the inter-arrival time threshold is raised, thereby allowing more content to be admitted to the SSD cache.

The activity of a SSD cache that is part of a CDN can vary significantly over time. For example, there may be a strong diurnal cycle, as caches in a particular location respond to changing demand over the duration of the day. Further, there may also be longer-time cycles, such as weekend vs. weekday demand variations. An important objective that a CDN is configured to meet is to serve more content at peak periods, when demand is at its highest. An admission control policy that is too short-term allows too little admission of content during peak periods, and allows too much admission of content during trough periods.

The admission control policy is therefore adjusted based on such periods. For example, content popularity metrics are computed over a demand cycle (e.g., daily, weekly, etc.), such that higher content admission rates are permitted during peak periods. The popularity database is configured to store the history of content admission over multiple demand cycles. To mitigate any excessive oscillatory behavior of the threshold, the admission thresholds as well as the adjustments to the threshold are bounded per adjustment cycle.

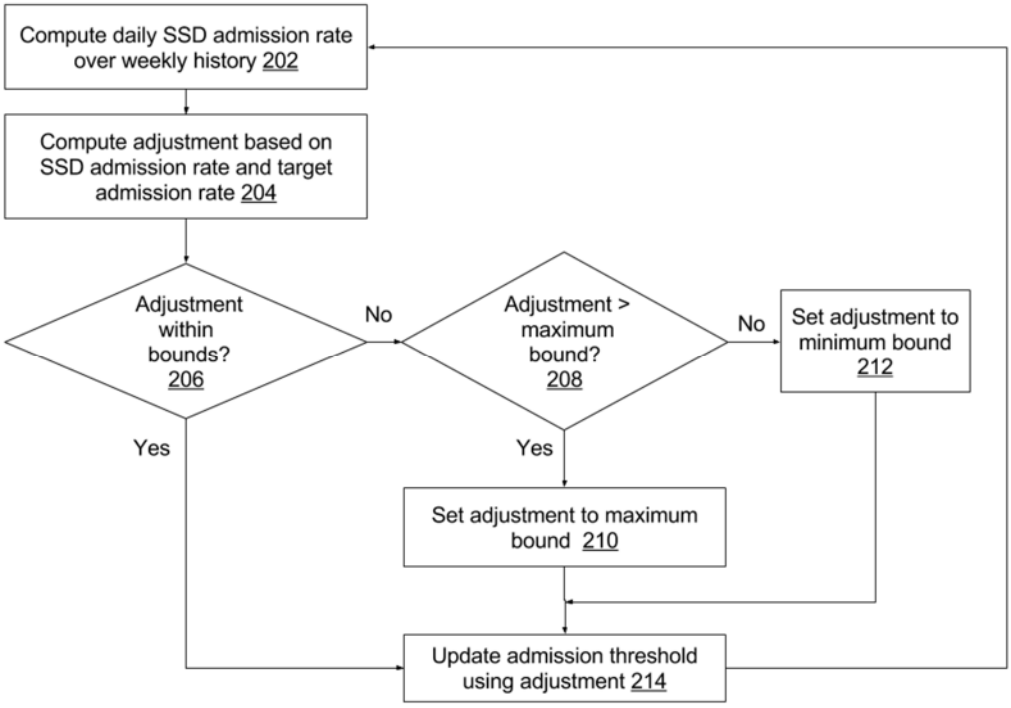

Fig. 2: Adjustment of admission control parameters

Fig. 2 illustrates a technique to adjust admission control parameters. The admission control threshold is adjusted at a predetermined interval (e.g., daily) using admission history over another predetermined interval (e.g., weekly). A target admission rate is computed based on the size of the SSD cache and SSD write parameters. The admission rate is computed (202), which gives an indication of the SSD device writes per day over the period for which history is maintained. The admission rate is compared to the target admission rate to generate an adjustment (204).

The adjustment is a multiplicative factor that adjusts the inter-arrival time to either raise or lower the popularity metric threshold. For example, a simple linear ratio of the desired and current admission (ingress) rate can be used. Once the adjustment is computed, it is checked against minimum and maximum bounds (206). This is done to prevent excessive oscillations of

the threshold. If the adjustment lies within the bounds, the computed adjustment is used to update the admission threshold (214). If the adjustment exceeds the maximum bound (208), the adjustment is limited to the allowed maximum bound (210). This value is used to update the admission threshold. If the adjustment falls below the minimum bound, the adjustment is set to the minimum bound (212). The admission threshold is accordingly updated (214).

The techniques take into account the write endurance properties of SSDs, as well as regular time-varying behavior, such as diurnal and weekly activity cycles of content popularity.

CONCLUSION

This disclosure provides techniques for admission control for writing content to SSD devices used as a cache tier for a CDN. Techniques provided herein allow for optimal utilization of SSD caches by dynamically adjusting admission thresholds such that the content writes to the devices do not exceed write endurance limits. The adjustments are controlled using bounds such that the content admission threshold remains within a stable range. The techniques provide improve lifetime of SSD devices used as caches in CDNs while maintaining content delivery performance.