

Diamonds are forever, files are not *

Surendar Chandra, Ashish Gehani and Xuwen Yu

University of Notre Dame

(surendar,ashish.gehani,xyu2)@nd.edu

This work was motivated by the observation that not all objects are equally important. In many scenarios, the importance also wanes with time. We leveraged this importance notions to reduce the management overhead for less important objects. In our system, all objects were explicitly annotated by the content creators using a temporal importance metric. The storage system used these hints to manage the reliability (by reducing the replication levels for less important objects), fidelity (by transcoding objects), performance (by moving more important objects to faster storage locations), security (by encrypting important objects with stronger encryption) and persistence (by reducing the persistence guarantees for less important objects). Our initial focus was on policies for managing the object persistence.

Our experiments were motivated by a university wide lecture video capture and storage system. Multimedia objects are large; depending on the quality used, compressed media objects can consume storage at the rate of multiple gigabytes per hour. Multimedia objects can eventually fill even the largest storage. Users were forced to either keep adding more storage on a regular basis or manually delete some objects. Keeping track of objects for deletion is cumbersome; many objects are orphaned and much of the storage remained under utilized.

The difficulty was that traditional storage systems considered that all objects were equally important. When the storage was full, the storage system denied requests to store newer objects even though old objects that were already stored in the system were potentially less important than the incoming object. Storage systems relied upon user intervention to reclaim the space consumed by less important objects. We relegated object reclamation policies to the storage system itself. This delegation notions is not new; Roscoe et. al. [1] described a system which evicted objects in a FIFO basis; the persistence achieved depends on the future object arrival rates. Zadok et. al. [2] defined two classes of objects wherein objects in the elastic quota class can be reclaimed as necessary. In their position paper, Douglass et. al. [3] introduced the notions of retention value of objects. Our analysis shows that, FIFO cache like replacement policies are agnostic towards object impor-

tance differences and do not provide adequate feedback on the overall state of the available storage to the user. Our goal is to investigate richer per object reclamation policies that are easier to use with predictable expected object lifetimes.

We introduce object *temporal importance* as a first class attribute to describe the expected longevity and importance of an object. *Lifetimes* are a monotonically decreasing importance function. *Importance* is a scalar metric; objects of higher *importance* can preempt objects of lower *importance*. The storage appears full for less *important* objects. The success of our system depends on the ability of the content creator to specify the required lifetimes. The requested lifetime depends on the content creator's understanding of the projected importance of the object as well as the expected space competition from other users who may use the same storage for different purposes. On the other hand, requesting conservatively longer durations is wasteful. We are investigating mechanisms that will provide enough feedback to the content creators in order to strike the right balance in choosing the object temporal importance.

We are exploring several temporal lifetime functions for their suitability in our target lecture capture system. We are conducting experiments to investigate the expressiveness, predictiveness and scalability of the lifetime functions. Our preliminary simulation results are encouraging. Currently, we are building a prototype to validate our simulations.

References

- [1] T. Roscoe and S. Hand, "Palimpsest: Soft-capacity storage for planetary-scale services," in *HotOS*, Kauai, HI, May 2003, pp. 127–132.
- [2] E. Zadok, J. Osborn, A. Shater, C. Wright, K.-K. Muniswamy-Reddy, and J. Nieh, "Reducing Storage Management Costs via Informed User-Based Policies," in *Mass Storage Systems and Technologies*, College Park, MD, April 2004, pp. 193–197.
- [3] F. Douglass, J. Palmer, E. S. Richards, D. Tao, W. H. Tetzlaff, J. M. Tracey, and J. Yin, "Position: short object lifetimes require a delete-optimized storage system," in *ACM SIGOPS European workshop: beyond the PC*. New York, NY, USA: ACM Press, Sept. 2004, p. 6.

*Supported in part by the U.S. National Science Foundation (IIS-0515674 and CNS-0447671)