

# Sprout: A functional caching approach to minimize service latency in erasure-coded storage

Vaneet Aggarwal<sup>1</sup>, Yih-Farn R Chen<sup>2</sup>, Tian Lan<sup>3</sup>, and Yu Xiang<sup>2</sup>

<sup>1</sup> School of IE, Purdue University, West Lafayette, IN 47907, USA

<sup>2</sup> AT&T Labs-Research, Bedminster, NJ 07921, USA

<sup>3</sup> Department of ECE, George Washington University, DC 20052, USA

vaneet@purdue.edu, chen@research.att.com, tlan@gwu.edu, yxiang@research.att.com

**Abstract**—Modern distributed storage systems often use erasure codes to protect against disk and node failures to increase reliability, while trying to meet the latency requirements of the applications and clients. Storage systems may have caches at the proxy or client ends in order to reduce the latency. In this paper, we consider a novel caching framework with erasure code called *functional caching*. Functional Caching involves using erasure-coded chunks in the cache such that the code formed by the chunks in storage nodes and cache combined are maximal-distance-separable (MDS) erasure codes. Based on the arrival rates of different files, placement of file chunks on the servers, and service time distribution of storage servers, an optimal functional caching placement and the access probabilities of the file request from different disks are considered. The proposed algorithm gives significant latency improvement in both simulations and a prototyped solution in an open-source, cloud storage deployment.

## I. INTRODUCTION

Erasure coding has seen itself quickly emerged as a promising technique to reduce the storage cost for a given reliability as compared to fully-replicated systems [2, 3]. It has been widely adopted in modern storage systems by companies like Facebook [4], Microsoft [5] and Google [6]. In these systems, the rapid growth of data traffic such as those generated by online video streaming, Big Data analytics, social networking and E-commerce activities has put a significant burden on the underlying networks of datacenter storage systems. Many researchers have begun to focus on latency analysis in erasure coded storage systems [7–12] and to investigate algorithms for joint latency optimization and resource management [12–14].

Historically, a key solution to relieve this traffic burden is caching [15]. By storing chunks of popular data at different locations closer to end-users, caching can greatly reduce congestion in the network and improve service delay for processing file requests. For example, Figure 1 shows a typical video storage architecture with video proxies and multiple video clients. It is very common for 20% of the video content to be accessed 80% of the time, so caching popular content at proxies significantly reduces the overall latency on the client side.

However, caching for datacenters where the files are encoded with erasure codes gives rise to new challenges. The

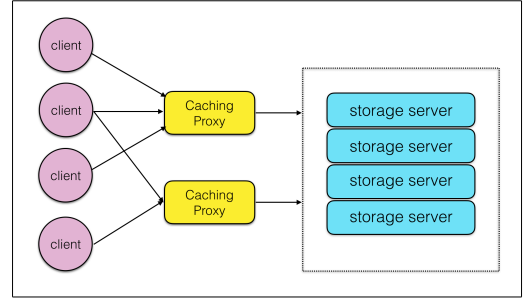


Fig. 1. Caching proxies for video storage to reduce the latency of client accesses for popular Video-on-Demand content

current results fall short of addressing the impact of erasure coding on latency and thus fail to providing insights on the optimal caching policy. First, using an  $(n, k)$  maximum-distance-separable (MDS) erasure code, a file is encoded into  $n$  chunks and can be recovered from any subset of  $k$  distinct chunks. Thus, file access latency in such a system is determined by the delay to access file chunks on hot storage nodes with the slowest performance. Significant latency reduction can be achieved by caching only a few hot chunks of each file (and therefore alleviating system performance bottlenecks), whereas caching additional chunks or even complete files (e.g., [16–19]) only has diminishing benefits. It is an open problem to design a caching policy that optimally apportions limited cache capacity among all files in an erasure coded storage to minimize overall access latency.

More importantly, caching the most popular data chunks is often optimal because the cache-miss rate and the resulting network load are proportional to each other. However, this may not be true for an erasure-coded storage, where cached chunks need not be identical to the chunks transferred from storage nodes. More precisely, leveraging the existing erasure coding, a function of the data chunks can be computed and cached, so that the constructed new chunks (i.e.,  $d$ ), along with the existing chunks, satisfy the property of being a new MDS code. It effectively expands the existing  $(n, k)$  code to an  $(n + d, k)$ , which leads to much lower access latency [12].

In this paper, we propose a new *functional caching* approach called *Sprout* that can efficiently capitalize on existing file coding in erasure-coded storage systems. In contrast to exact caching that stores  $d$  chunks identical to original copies, our functional caching approach forms  $d$  new data chunks, which together with the existing  $n$  chunks satisfy the property of

being an  $(n + d, k)$  MDS code. Thus, the file can now be recovered from any  $k$  out of  $n + d$  chunks (rather than  $k$  out of  $n$  under exact caching), effectively extending coding redundancy, as well as system diversity for scheduling file access requests. The proposed functional caching approach saves latency due to more flexibility to obtain  $k - d$  chunks from the storage system at a very minimal additional computational cost of creating the coded cached chunks. To the best of our knowledge, this is the first work studying functional caching for erasure-coded storage and proposing an analytical framework for cache optimization.

Most of the caching strategies in today's data centers cache complete files [20–22]. This flexibility without erasure coding has been recently explored in [23]. Partial number of chunks in cache give more flexibility as compared to caching the entire file. This is because the connection from some servers may be better than others and caching a smaller number of chunks can aid avoiding the bottleneck. The proposed optimization includes as a special case complete file caching and thus the performance can be no worse than caching the entire file in the cache. This paper gives a novel approach, functional caching, which is an efficient approach to cache partial chunks since any  $k - d$  of the  $n$  servers can be used thus increasing flexibility as compared to using any  $k - d$  of the remaining  $n - d$  servers when  $d$  exact chunks are copied in the cache. This additional flexibility shows that the latency with functional caching is no higher than the strategy where part of the chunks on the servers are cached as such.

While quantifying service latency in erasure-coded storage systems is an open problem, we generalize previous results on probabilistic scheduling policy [11, 12] that distributes the file requests to cache and storage nodes with optimized probabilities, and derive a closed-form upper bound on mean service latency for the proposed functional caching approach. The latency bound is obtained using order-statistic [24] analysis and it works on erasure-coded storage systems with arbitrary cache size and data chunk placement.

This analytical latency model for functional caching enables us to formulate a cache-content optimization problem. This problem is an integer optimization problem, which is very difficult to solve. Towards this end, for given data chunk placement and file request arrival rates, we propose a heuristic algorithm that iteratively identifies files whose service latency benefits most from caching and constructs new functional data chunks until the cache is filled up. The algorithm can be efficiently computed to allow online cache optimization and management with time-varying arrival rates.

The proposed algorithm is an iterative algorithm, which converges within a few iterations in our conducted experiments and it was validated by the numerical results. For 1000 files, we find that the algorithm converges within 20 iterations. The file latency decreases as a convex function as the cache size increases thus showing diminishing returns for the increase in cache size. We also find that it is suboptimal in general to have all  $k$  chunks of an  $(n, k)$  coded file in the cache. Further, the cache placement depends heavily on the file arrival rates, storage server service time distributions, as well as the content placement on the files. If a high arrival-rate file is placed on servers which have a less overall load, this file may not have

any contribution in the cache. Thus, the proposed algorithm accounts for all the system parameters to effectively find the cache placement. The proposed algorithm is prototyped using *Ceph*, an open-source erasure-coded storage system [25] and tested on a real-world storage testbed with an emulation of real storage workload. We find that caching helps improve the latency significantly.

The key contributions of our paper include:

- We propose a new functional caching scheme that leverages existing file coding in erasure-coded storage systems, and quantify its service latency through probabilistic request scheduling.
- Based on the service latency analysis, an optimization problem is formulated which minimizes the average latency of the files. This problem has integer constraints due to the integer number of chunks in the cache.
- An iterative algorithm is developed to optimize cache content. The proposed algorithm takes file placement, service time distributions, and file arrival rates into account to find the cache placement which optimizes the service latency.
- The simulation results show that the proposed algorithm converges within a few iterations.
- The prototype implementation of the functional caching scheme and the cache optimization algorithm using *Ceph* are used to validate significant latency reduction on a real-world storage testbed. As compared to the *Ceph*'s LRU (least recently used) caching algorithm, the algorithm proposed in this paper reduces latency by 24.93% on an average for all tested workloads in the prototype implementation.

The remainder of this paper is organized as follows. Section II provides related work for this paper. In Section III, we describe the system model used in the paper with a description of functional caching. Section IV formulates the cache optimization problem and develops an iterative algorithmic solution. Prototype and evaluation results are included in Section V. Section VI concludes the paper.

## II. RELATED WORK

Quantifying exact latency for an erasure-coded storage system is an open problem. Recently, there has been a number of attempts at finding latency bounds for an erasure-coded storage system [7–14, 26–28]. In this paper, we utilize the probabilistic scheduling policy developed in [11, 12] and analyze the impact of caching on the service latency of erasure-coded storage. Even though exact characterization of latency is open, probabilistic scheduling has been shown to be optimal for latency tail index, when the file size is Pareto distributed [29].

Caches are a critical resource in data centers; however, there is little work on caching for erasure coded distributed storage. The problem of online cache management (i.e., decisions for evicting a data chunk currently in the cache to make room for a new chunk) has been studied for networks with distributed caches [30, 31]. Cache replacement strategy called LRU (Least Recently Used) is widely used in managing buffer cache due to its simplicity [32–36]. Recently, a steady-state characterization

of various cache policies is developed in [37], and new coded caching schemes to enable multicast opportunities for bandwidth minimization are proposed in [15, 38]. Recently, caching in erasure-coded storage has been studied [39] where the cache stores the files in their entirety. In contrast, this paper allows for partial chunks, which are functionally different from those stored on the servers.

Much of the work on caching in data centers is focused on specialized application caches, such as Facebook’s photo-serving stack [40], Facebook’s social graph store [41], memcached [42], or explicit cloud caching services [43, 44]. However, in the presence of coding, new challenges arise. First, the chunks are distributed over multiple servers and a part of the chunks can be in the cache. Thus, it is not necessary for the complete file to be in the cache. Second, the latency calculation for a file depends on the placement of the files and the request of the files from  $k$  out of  $n$  servers. In this paper, we deal with these challenges to consider a novel caching strategy. To the best of our knowledge, this is the first work on caching with erasure coded files on distributed storage servers accounting for the latency in file retrievals, based on the estimated arrival rates.

Coded caching for a single server with multi-cast link to different users has been considered in [45–48]. This does not account for multiple distributed storage servers and latency to get the content. An extension of the approach to distributed storage systems is considered recently in [49], where multiple cache-enabled clients connected to multiple servers through an intermediate network. However, the impact of coding on the servers, and limited service rate of different servers is not taken into account. The key idea in this set of works uses a coded version of different files in the cache which helps in the case when users request different files with certain probabilities. The gain of the approach is due to the model where a message from the server can be received at multiple nodes and thus combined with coded content in the cache, one chunk from the server can help give a chunk for different files at different clients. In this paper, we do not have multicast links to different clients and thus coding across files in the cache is not used.

The functional repair was introduced in [3] for repairing a failed chunk with a new chunk such that the storage code satisfies the same properties even after repair. Thus, the replaced content can be different. In this paper, we use functional caching to have chunks in the cache such that the file can be recovered from any of  $k$  chunks from a combination of disk and cache contents where  $(n, k)$  code is used for the files in the storage system.

### III. SYSTEM MODEL

We consider a distributed storage system consisting of  $m$  heterogeneous storage nodes, denoted by  $\mathcal{M} = \{1, 2, \dots, m\}$ . To distributively store a set of  $r$  files, indexed by  $i = 1, \dots, r$ , we partition each file  $i$  into  $k_i$  fixed-size chunks<sup>1</sup> and then encode it using an  $(n_i, k_i)$  MDS erasure code to generate  $n_i$

distinct chunks of the same size for file  $i$ . The encoded chunks are stored on the disks of  $n_i$  distinct storage nodes. A set  $\mathcal{S}_i$  of storage nodes, satisfying  $\mathcal{S}_i \subseteq \mathcal{M}$  and  $n_i = |\mathcal{S}_i|$  is used to store file  $i$ . Therefore, each chunk is placed on a different node to provide high reliability in the event of node or network failures. The use of  $(n_i, k_i)$  MDS erasure code allows the file to be reconstructed from any subset of  $k_i$ -out-of- $n_i$  chunks, whereas it also introduces a redundancy factor of  $n_i/k_i$ .

The files are accessed (to be read in their entirety) by compute servers located in the same datacenter. A networked cache of size  $C$  is available at each compute server to store a limited number of chunks of the  $r$  files in its cache memory. File access requests are modeled by a non-homogenous Poisson process. We make the assumption of time-scale separation, such that system service time is divided into multiple bins, each with different request arrival rates, while the arrival rates within each bin remain stationary. This model allows us to analyze cache placement and latency in steady-state for each time bin, and by varying arrival rates for different time bins, to also take into account time-varying service rates during busy and off-peak hours. Let  $\lambda_{i,j,t}$  be the arrival rate of file- $i$  requests at compute server  $j$  in time bin  $t$ . Since each cache serves a single compute server, we consider a separate optimization for each cache and suppress server index  $j$  in the notations. Let  $d_i \leq k_i$  (chunks) be the size of cache memory allocated to storing file  $i$  chunks. These chunks in the cache memory can be both prefetched in an offline fashion during a placement phase [15] (during hours of low workload) and updated on the fly when a file  $i$  request is processed by the system.

**Functional Caching.** Under functional caching,  $d_i$  new coded data chunks of file  $i$  are constructed and cached, so that along with the existing  $n_i$  chunks satisfy the property of being an  $(n_i + d_i, k_i)$  MDS code. In this paper, we use Reed-Solomon codes [53] to generate MDS codes for arbitrary  $(n_i + k_i, k_i)$ . Thus a subset of the encoded chunks generated using these codes satisfy the desired property for functional caching. More precisely, for given erasure coding and chunk placement on storage nodes and cache, a request to access file  $i$  can be processed using  $d_i$  cached chunks in conjunction with  $k_i - d_i$  chunks on distinct storage nodes. After each file request arrives at the storage system, we model this by treating the file request as a batch of  $k_i - d_i$  chunk requests that are forwarded to appropriate storage nodes, as well as  $d_i$  chunk requests that are processed by the cache. Each storage node buffers requests in a common queue of infinite capacity and process them in a FIFO manner. The file request is served when all  $k_i$  chunk requests are processed. Further, we consider chunk service time  $\mathbf{X}_j$  of node  $j$  with arbitrary distributions, whose statistics can be inferred from existing work on network delay [54, 55] and file-size distribution [56, 57].

**An Illustrative Example.** Consider a datacenter storing a single file using a  $(6, 5)$  MDS code. The file is split into  $k_i = 5$  chunks, denoted by  $A_1, A_2, A_3, A_4, A_5$ , and then linearly encoded to generate  $n_i = 6$  coded chunks  $F_1 = A_1, F_2 = A_2, F_3 = A_3, F_4 = A_4, F_5 = A_5$  and  $F_6 = A_1 + A_2 + A_3 + A_4 + A_5$  in a finite field of order at-least 5. In this example, we compare the latency of 2 different cache policies: (i) Exact caching in Rack 3 that stores chunks

<sup>1</sup>While we make the assumption of fixed chunk size here to simplify the problem formulation, all results in this paper can be easily extended to variable chunk sizes. Nevertheless, fixed chunk sizes are indeed used by many existing storage systems [50–52].

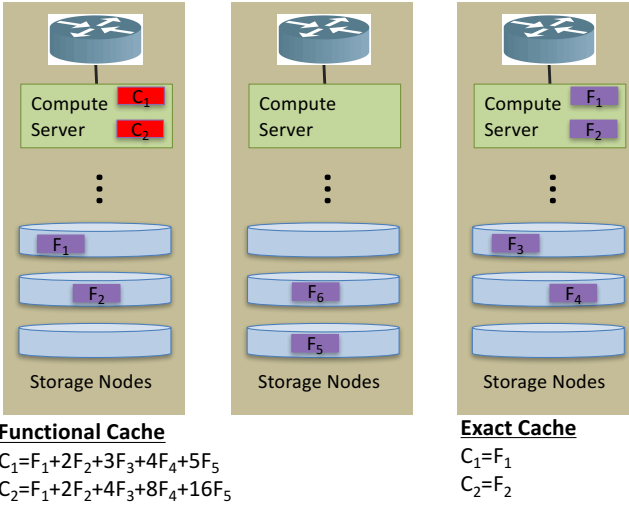


Fig. 2. An illustration of functional caching and exact caching in an erasure-coded storage system with one file using a  $(6, 5)$  erasure code. The file is split into  $k_i = 5$  chunks, denoted by  $A_1, A_2, A_3, A_4, A_5$ , and then linearly encoded to generate  $n_i = 6$  coded chunks  $F_1 = A_1, F_2 = A_2, F_3 = A_3, F_4 = A_4, F_5 = A_5$ , and  $F_6 = A_1 + A_2 + A_3 + A_4 + A_5$ .

$F_1$  and  $F_2$ , and (ii) Functional caching in Rack 1 that stores 2 coded chunks  $C_1$  and  $C_2$ .

Due to the use of  $(6, 5)$  MDS code, the file can be reconstructed using any 5 out of 6 chunks. Two compute servers in the datacenter access this file and each is equipped with a cache of size  $C = 2$  chunks as depicted in Figure 2. The compute server on the right employs an exact caching scheme and stores chunks  $F_1, F_2$  in the cache memory. Thus, 3 out of 4 remaining chunks (i.e.,  $F_3, F_4, F_5$  or  $F_6$ ) must be retrieved to access the file. Chunks  $F_1, F_2$  and thus their host nodes will not be selected for processing the requests even if the servers in Rack 1 are least congested. Since file access latency is determined by the time required to retrieve all 5 selected chunks, caching  $F_1, F_2$  may not necessarily reduce file access latency if other chunks in Rack 2 and 3 are currently the performance bottleneck.

We show that lower latency can always be achieved in functional caching by exploiting the erasure code structure. More specifically, the compute server on the left generates  $d_i = 2$  new coded chunks, i.e.,  $C_1 = A_1 + 2A_2 + 3A_3 + 4A_4 + 5A_5$  and  $C_2 = A_1 + 2A_2 + 4A_3 + 8A_4 + 16A_5$ , and saves them in its cache memory. It is easy to see that chunks  $F_1, \dots, F_6$  and  $C_1, C_2$  now form a  $(8, 5)$  erasure code. Thus, the file can be retrieved by accessing  $C_1, C_2$  in the cache together with any 3 out of 6 chunks from  $F_1, \dots, F_6$ . This allows an optimal request scheduling mechanism to select the least busy chunks/nodes among all 5 possible candidates in the system so that the service latency is determined by the best 3 storage nodes with minimum queuing delay. In other words, the coded chunks  $C_1, C_2$  can always be used to replace the slowest chunks required under exact caching policy, resulting in smaller latency. Consider file access latency, which is determined by the time required to access any 5 chunks on the storage servers. Under an exact caching of file copies  $F_1$  and  $F_2$ , any request to access the file must select 3 other files chunks hosted by storage servers. Thus, access latency in this case is determined by the minimum time to retrieve any

3 chunks out of  $F_3, F_4, F_5$  and  $F_6$ . In contrast, under our proposed functional caching policy, given the availability of coded chunks  $C_1$  and  $C_2$  in cache, any 2 chunks from  $F_1, F_2, F_3, F_4, F_5$  and  $F_6$  is sufficient for recovery, which leads to smaller latency if nodes storing  $F_1$  and  $F_2$  are less congested and faster to access. This approach effectively extends to an  $(8, 5)$  erasure code and guarantees lower access latency and also enables better load balancing due to a higher degree of flexibility in request scheduling and chunk selection.

In order to have an  $(n, k)$  coded file in the storage server, we can construct chunks by using an  $(n + k, k)$  MDS code, where  $n$  chunks are stored in the storage server. The remaining  $k$  out of the  $n + k$  coded chunks are assigned to be in part in the cache based on the contents of the file in the cache. Thus, irrespective of the value of  $d \leq k$ , we ascertain that  $(n + d, k)$  code, formed with  $n$  coded chunks in the storage server and  $d$  coded chunks in the cache will be MDS.

In the proposed system, the placement in each time-bin is decided based on the predicted arrival rates in the time bin. The time bin can either be a fixed time or dynamic based on a significant change of the predicted arrival rates. At the start of the time-bin, new cache placement is found using the optimized algorithm. For each file that has a lower number of chunks in the new time bin, the decreased contents are removed from the cache. For the files for which the cache contents increase in a time bin, we wait for the file to be accessed. When the file is accessed, the file contents are gathered and the required new chunks are generated to be placed in the cache. Thus, the change of cache content does not place any additional network overhead and the cache contents of a file are added only when it is first accessed in the new time bin. This process can be further improved improving latency till convergence to the new cache content in the new time bin by not letting all the chunks which have to be removed all simultaneously but removing as needed based on the added chunks.

#### IV. OPTIMIZED DISTRIBUTED STORAGE WITH CACHE

In this section, we quantify the mean service latency for file access with functional caching. The result enables us to formulate a cache optimization problem for latency minimization and develop an efficient algorithm solution.

##### A. Formulation of Optimization

At time  $t$ , we consider the cache optimization problem, which decides the optimal number  $d_{i,t}$  of file- $i$  chunks to store in the cache memory, satisfying cache capacity constraint  $\sum_{i=1}^r d_{i,t} \leq C$ , in order to minimize mean service latency of all files. Under functional caching, each file- $i$  request is served by accessing  $d_{i,t}$  chunks in the cache, along with  $k_i - d_{i,t}$  distinct chunks that are selected from  $n_i$  storage nodes. Thus, the latency to access file  $i$  under functional caching is determined by the maximum processing (queuing) delay of the selected  $k_i - d_{i,t}$  storage nodes. Quantifying service latency in such erasure-coded system is an open problem. In this paper, we use probabilistic scheduling proposed in [12] to derive an upper bound on the average file latency.



The key idea is to forward each file- $i$  request to a set of  $k_i - d_{i,t}$  storage nodes (denoted by  $\mathcal{A}_{i,t} \subseteq \mathcal{S}_i$ ) with some predetermined probabilities  $\{\pi_{i,j,t} \in [0, 1], \forall i, j, t\}$  for  $j \in \mathcal{A}_{i,t}$ . Each node then manages a local queue and process chunk requests with service rate  $\mu_j$ . While the local queues are not independent due to coupled request arrivals, we can leverage order statistic analysis to derive an upper bound of mean service latency in closed-form [12]. The result is then optimized over probabilities  $\pi_{i,j,t}$  to obtain the tightest bound. Let  $\mathbf{Q}_{j,t}$  be the (random) waiting time a chunk request spends in the queue of node  $j$  in time-bin  $t$ . Using the functional caching approach, requests of file  $i$  see mean latency  $\bar{T}_{i,t}$  given by

$$\bar{T}_{i,t} = \mathbb{E} \left[ \mathbb{E}_{\mathcal{A}_{i,t}} \left( \max_{j \in \mathcal{A}_{i,t}} \{\mathbf{Q}_{j,t}\} \right) \right], \quad (1)$$

where the first expectation is taken over system queuing dynamics and the second expectation is taken over random dispatch decisions  $\mathcal{A}_{i,t}$ . We note that queuing delay  $\mathbf{Q}_{j,t}$ 's are dependent due to coupled request arrivals. Therefore, an exact queuing analysis of the system is intractable. We use the technique in [12] to derive an upper bound of (1).

We denote  $\mathbf{X}_j$  as the service time per chunk at node  $j$ , which has an arbitrary distribution satisfying finite mean  $\mathbb{E}[\mathbf{X}_j] = 1/\mu_j$ , variance  $\mathbb{E}[\mathbf{X}_j^2] - \mathbb{E}[\mathbf{X}_j]^2 = \sigma_j^2$ , second moment  $\mathbb{E}[\mathbf{X}_j^2] = \Gamma_j^2$ , and third moment  $\mathbb{E}[\mathbf{X}_j^3] = \hat{\Gamma}_j^3$ . These statistics can be readily inferred from existing work on network delay [54, 55] and file-size distribution [56, 57]. Following [12], an upper bound on the expected latency is given as follows.

**Lemma 1.** *The expected latency  $\bar{T}_{i,t}$  of file  $i$  in time-bin  $t$  under probabilistic scheduling is upper bounded by  $\bar{U}_{i,t}$ , given by*

$$\bar{U}_{i,t} = \min_{z_{i,t} \in \mathbb{R}} \left\{ z_{i,t} + \sum_{j \in \mathcal{S}_{i,t}} \frac{\pi_{i,j,t}}{2} (\mathbb{E}[\mathbf{Q}_{j,t}] - z_{i,t}) + \sum_{j \in \mathcal{S}_{i,t}} \frac{\pi_{i,j,t}}{2} \left[ \sqrt{(\mathbb{E}[\mathbf{Q}_{j,t}] - z_{i,t})^2 + \text{Var}[\mathbf{Q}_{j,t}]} \right] \right\} \quad (2)$$

where

$$\mathbb{E}[\mathbf{Q}_{j,t}] = \frac{1}{\mu_j} + \frac{\Lambda_{j,t} \Gamma_j^2}{2(1 - \rho_{j,t})}, \quad (3)$$

$$\text{Var}[\mathbf{Q}_{j,t}] = \sigma_j^2 + \frac{\Lambda_{j,t} \hat{\Gamma}_j^3}{3(1 - \rho_{j,t})} + \frac{\Lambda_{j,t}^2 \Gamma_j^4}{4(1 - \rho_{j,t})^2}, \quad (4)$$

where  $\rho_{j,t} = \Lambda_{j,t}/\mu_j$  is the request intensity at node  $j$ , and  $\Lambda_{j,t} = \sum_{i=1}^r \lambda_{i,t} \pi_{i,j,t}$  is the mean arrival rate at node  $j$ . The bound is tight in the sense that there exists a distribution of  $\mathbf{Q}_{j,t}$  such that (2) is satisfied with exact equality.

*Proof.* The proof follows on the same lines as in [12]. A brief detail into the proof idea is provided in Appendix A (where the time-index  $t$  is omitted) to make the paper more self-contained.  $\square$

**Remark 1.** *The upper bound of Lemma 1 can be improved if higher ordered moment statistics of service time are known, though we do not consider them in this work due to tractability in latency optimization. This bound, even without the additional moments, has been shown to outperform existing bounds and is indeed very close the actual latency [12].*

Under the assumption of time-scale separation, each time bin has stationary request arrival rates and is long enough for cache content to reach a steady state. We consider the cache optimization in a single time bin to minimize weighted mean latency quantified by Lemma 1. This optimization will be performed repeatedly at the beginning of each time bin to address time-varying request arrival rates during busy and off-peak hours.

We now formulate the cache optimization in a single time-bin. The optimization is over cache content placement  $d_{i,t}$ , scheduling probabilities  $\pi_{i,j,t}$ , and auxiliary variable  $z_{i,t}$  in the upper bound. Let  $\hat{\lambda}_t = \sum_{i=1}^r \lambda_{i,t}$  be the total arrival rate, so  $\lambda_{i,t}/\hat{\lambda}_t$  is the fraction of file  $i$  requests, and average latency of all files is given by  $\sum_{i=1}^r (\lambda_{i,t}/\hat{\lambda}_t) \bar{T}_{i,t}$ . Our objective is to minimize an average latency objective, i.e.,

$$\begin{aligned} \min \quad & \sum_{i=1}^r \frac{\lambda_{i,t}}{\hat{\lambda}_t} \bar{U}_{i,t} \\ \text{s.t.} \quad & (2), (3), (4), \sum_{j=1}^m \pi_{i,j,t} = k_i - d_{i,t}, \pi_{i,j,t}, d_{i,t} \geq 0, \\ & \sum_{i=1}^r d_{i,t} \leq C, \pi_{i,j,t} = 0 \text{ for } j \notin \mathcal{S}_i, \pi_{i,j,t} \leq 1, \\ & z_{i,t} \geq 0, d_{i,t} \in \mathbb{Z}. \\ \text{var.} \quad & \pi_{i,j,t}, d_{i,t}, z_{i,t} \quad \forall i, j, t. \end{aligned} \quad (5)$$

Here the constraints  $\sum_{j=1}^m \pi_{i,j,t} = k_i - d_{i,t}$  and  $\pi_{i,j,t} \leq 1$  ensure that  $k_i - d_{i,t}$  distinct storage nodes (along with  $d_{i,t}$  chunks in the cache) are selected to process each file request, following probabilistic scheduling in [12]. Clearly, storage nodes without desired chunks cannot be selected, i.e.,  $\pi_{i,j,t} = 0$  for  $j \notin \mathcal{S}_i$ . Finally, the cache has a capacity constraint  $\sum_{i=1}^r d_{i,t} \leq C$ .

Solving the cache optimization gives us the optimal cache content placement and scheduling policy to minimize file access latency. We note that the constraint  $z_{i,t} \geq 0$  is not needed if none of the files is completely in the cache. However, the latency bound does not hold if the file is completely in the cache since in that case the bound is  $z_{i,t}$  in the above expression. In order to avoid having indicators representing the constraint on  $z_{i,t} = 0$  if the file is in the cache, we only consider  $z_{i,t} \geq 0$  making the latency bound hold irrespective of the number of chunks in the cache. This problem can be rewritten as follows.

### Distributed Storage with Caching:

$$\min \sum_{i=1}^r \lambda_{i,t} z_{i,t} / \hat{\lambda}_t + \sum_{i=1}^r \sum_{j=1}^m \frac{\lambda_{i,t} \pi_{i,j,t}}{2\hat{\lambda}_t} \left[ X_{i,j,t} + \sqrt{X_{i,j,t}^2 + Y_{j,t}} \right] \quad (6)$$

$$s.t. \quad X_{i,j,t} = \frac{1}{\mu_j} + \frac{\Lambda_{j,t} \Gamma_j^2}{2(1 - \rho_{j,t})} - z_{i,t}, \quad \forall j \quad (7)$$

$$Y_{j,t} = \sigma_j^2 + \frac{\Lambda_{j,t} \hat{\Gamma}_j^3}{3(1 - \rho_{j,t})} + \frac{\Lambda_{j,t}^2 \Gamma_j^4}{4(1 - \rho_{j,t})^2}, \quad \forall j \quad (8)$$

$$\rho_{j,t} = \Lambda_{j,t} / \mu_j < 1, \quad \Lambda_{j,t} = \sum_{i=1}^r \pi_{i,j,t} \lambda_{i,t} \quad \forall j \quad (9)$$

$$\sum_{j=1}^m \pi_{i,j,t} = k_i - d_{i,t}, \quad \pi_{i,j,t} \in [0, 1], \quad z_{i,t} \geq 0, \quad (10)$$

$$\pi_{i,j,t} = 0 \quad \forall j \notin \mathcal{S}_i, \quad \sum_{i=1}^r d_{i,t} \leq C, \quad d_{i,t} \in \mathbb{Z}^+ \quad (11)$$

$$\text{var.} \quad z_{i,t}, \quad d_{i,t}, \quad \pi_{i,j,t}, \quad \forall i, j.$$

### B. Proposed Algorithm

The proposed cache optimization problem in (6)-(11) is an integer optimization problem, since the number  $d_{i,t}$  of functional chunks in the cache must be integers. To solve this problem, we propose a heuristic algorithm, which iteratively identifies the files that benefit most from caching, and constructs/stores functional chunks into cache memory accordingly. We first note that the variable  $d_{i,t}$  can be absorbed into scheduling decision  $\pi_{i,j,t}$  because of the equality constraint  $d_{i,t} = k_i - \sum_{j=1}^m \pi_{i,j,t}$ . Thus, there are two set of variables -  $z_{i,t}$ , and  $\pi_{i,j,t}$  - we need to consider. It is easy to show that the objective function is convex in both these variables, however there is an integer constraint on  $\sum_{j=1}^m \pi_{i,j,t}$  due to the integer requirement of  $d_{i,t}$ .

The algorithm employs an alternating minimization over two dimensions - the first decides on  $z_{i,t}$  given  $\pi_{i,j,t}$ , and the second decides on  $\pi_{i,j,t}$  given  $z_{i,t}$ . The first problem is convex, and can be easily solved by gradient descent. However, the second problem has integer constraint. In order to deal with this, we first remove integer constraint to solve the problem. Then, a certain percentage of files whose fractional part of content accessed from the disk is highest are added a part in the disk to make the part in the disk as integers. The optimization over  $\pi_{i,j,t}$  keeps running until  $\sum_{j=1}^m \pi_{i,j,t}$  for all files is an integer. In particular, we derive the two sub-problems that need to be solved as follows.

We define the problem *Prob\_Z* for given  $\pi_{i,j,t}$  as

$$\min \sum_{i=1}^r \lambda_{i,t} z_{i,t} / \hat{\lambda}_t + \sum_{i=1}^r \sum_{j=1}^m \frac{\lambda_{i,t} \pi_{i,j,t}}{2\hat{\lambda}_t} \left[ X_{i,j,t} + \sqrt{X_{i,j,t}^2 + Y_{j,t}} \right] \quad (12)$$

$$s.t. \quad (7), (8), (9), \quad z_{i,t} \geq 0$$

$$\text{var.} \quad z_{i,t}, \quad \forall i.$$

---

### Algorithm 1 Our Proposed Algorithm for Distributed Storage with Caching

---

Initialize  $c = 0$  and feasible  $(z_{i,t}, \pi_{i,j,t} \quad \forall i, j)$

Compute current objective value  $B^{(0)}$

Initialize  $c = 0$  and feasible  $(z_{i,t}, \pi_{i,j,t} \quad \forall i, j)$

Compute current objective value  $B^{(0)}$

**do**

Solve Convex Problem *Prob\_Z* to get  $z_{i,t}$  for given

$\pi_{i,j,t}$  for all  $i$

Set  $k_{L,i,t} = 0, k_{U,i,t} = k_i$

**do**

Solve Convex Problem *Prob\_II* to get  $\pi_{i,j,t}$  for given  $z_{i,t}, k_{L,i,t}, k_{U,i,t}$  for all  $i, j$

Let  $i_1 = \arg \max$  (fractional part of  $\sum_{j=1}^m \pi_{i,j,t}$ )

$k_{L,i_1,t} = k_{U,i_1,t} = \text{ceil}(\sum_{j=1}^m \pi_{i,j,t})$

**while**  $\sum_{i=1}^r \text{frac}(\sum_{j=1}^m \pi_{i,j,t}) > 0$

Compute new objective value  $B^{(c+1)}$ , Update  $c = c+1$

**while**  $B^{(c)} - B^{(c-1)} > \epsilon$

---

We define the problem *Prob\_II* for given  $z_{i,t}, k_{U,i,t}, k_{L,i,t}$  as

$$\min \sum_{i=1}^r \lambda_{i,t} z_{i,t} / \hat{\lambda}_t + \sum_{i=1}^r \sum_{j=1}^m \frac{\lambda_{i,t} \pi_{i,j,t}}{2\hat{\lambda}_t} \left[ X_{i,j,t} + \sqrt{X_{i,j,t}^2 + Y_{j,t}} \right] \quad (13)$$

$$s.t. \quad (7), (8), (9), \quad \pi_{i,j,t} \in [0, 1],$$

$$K_{L,i,t} \leq \sum_{j=1}^m \pi_{i,j,t} \leq k_{U,i,t}, \quad (14)$$

$$\pi_{i,j,t} = 0 \quad \forall j \notin \mathcal{S}_i, \quad \sum_{i=1}^r (k_i - \sum_{j=1}^m \pi_{i,j,t}) \leq C, \quad (15)$$

$$\text{var.} \quad \pi_{i,j,t}, \quad \forall i, j.$$

The problem *Prob\_Z* optimizes over  $z_{i,t}$  given  $\pi_{i,j,t}$ . This problem is convex with only one linear constraint  $z_{i,t} \geq 0$ . In order to solve this problem, we can use standard gradient descent, with making  $z_{i,t}$  as zero if the solution is negative in each iteration. The problem *Prob\_II* assumes that the number of total chunks of a file  $i$  accessed from the disk is between  $k_{L,i,t}$  and  $k_{U,i,t}$ . As we decide the number of chunks in the cache for each file, these two bounds will become equal. This problem is also convex, and can be solved using projected gradient descent. With algorithmic solution to these two sub-problems, the algorithm for Distributed Storage with Caching is given in Algorithm 1.

**Lemma 2.** *The 2 sub-problems, Prob\_Z and Prob\_II, are both convex.*

**Remark 2.** *The proposed algorithm 1 is guaranteed to converge. Since the inner loop that solves convex problem Prob\_II iteratively determines the placement of chunks in cache one-by-one, until the entire cache space is filled up, it runs at most  $r$  iterations. The outer loop that solves convex optimization Prob\_Z generates a monotonically decreasing sequence of objective values and is also guaranteed to converge within finite number of iterations.*

We note that the inner do-while logic to deal with integer optimization runs at most  $r$  times. Since  $r$  may be large, rather than choosing one index  $i_i$ , we choose a ceiling of certain fraction of file indices among those which have fractional content in the cache. This makes the loop run in  $O(\log r)$ . Thus, each outer loop runs  $O(\log r)$  convex problems. The algorithm will be solved repeatedly for each time bin to guide the update of cache content for service latency minimization.

## V. NUMERICAL RESULTS, IMPLEMENTATION AND EVALUATION

In this section, we evaluate our proposed algorithm for functional cache optimization, through both simulation and implementation in an open-source, distributed filesystem.

### A. Numerical Results Setup

We simulated our algorithm in a cluster of  $m = 12$  storage servers, holding  $r = 1000$  files of size 100 MB each using a (7,4) erasure code. Unless stated otherwise, cache size remains as 500 times of the chunk size (i.e., 500 times of 25 MB). The arrival rate for each file is set at a default value of  $\lambda_i = 0.000156/\text{sec}$ ,  $0.000156/\text{sec}$ ,  $0.000125/\text{sec}$ ,  $0.000167/\text{sec}$ ,  $0.000104/\text{sec}$  for every five out of the 1000 files of each size. It gives an aggregate arrival rate of all files to be  $0.1416/\text{sec}$ . The inverse of mean service times for the 12 servers are set based on measurements of real service time in a distributed storage system, which we obtained from a previous work [12], and they are  $\{0.1, 0.1, 0.1, 0.0909, 0.0909, 0.0667, 0.0667, 0.0769, 0.0769, 0.0588, 0.0588\}$  for the 12 storage servers respectively. The placement of files on the servers is chosen at random, unless explicitly specified. The major objective of our simulation is to validate our latency analysis of erasure-coded storage with caching in those areas that are hard to implement in a real test-bed or are difficult to measure in experiments.

### B. Numerical Results

**Convergence of Algorithm:** We implemented our cache optimization algorithm using MOSEK, a commercial optimization solver, to project the gradient descent solution to the feasible space for Prob\_II. For 12 distributed storage servers in our testbed, Figure 3 demonstrates the convergence of our algorithm in one time-bin, which optimizes the average latency of all files over request scheduling  $\pi_{i,j,t}$  and cache placement  $d_{i,t}$  for each file. We assume 1000 files, each using a (7,4) erasure code and of size 100 MB, divided evenly into five groups with the arrival rates of the five groups as mentioned above. The convergence of the proposed algorithm is depicted in Fig. 3 for cache size  $C \times 25$  MB. For  $C = 4000$ , four chunks of each file can be in the cache. A random initialization is chosen for  $C = 100$ , while the converged solution for  $C = 100$  is taken as initialization for  $C = 200$  and so on. We note that the algorithm converges within a few iterations, and converges in less than 20 iterations with a threshold of 0.01 on latency for all cache size values in Figure 3.

**Impact of Cache Size:** To validate that our proposed algorithm can effectively update file content placement in cache to optimize average latency in the storage system, we plot average latency for  $r = 1000$  files of size 100MB in

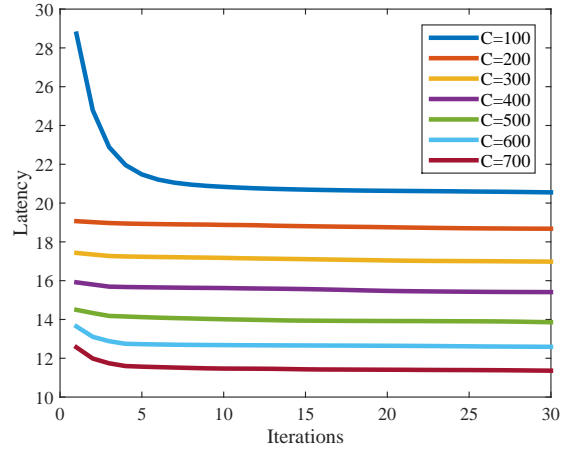


Fig. 3. Convergence of the proposed algorithm for a system with  $r = 1000$  files each of size 100MB and using a cache size of  $C \times 25$ M. The algorithm efficiently computes a solution in less than 20 iterations with a low tolerance  $\epsilon = 0.01$  for different cache sizes. A random initialization is chosen for  $C = 100$ , while the converged solution for  $C = 100$  is taken as initialization for  $C = 200$  and so on.

one time bin with designated request arrival rate for each file, while cache size varies from 0 to 4000 times of the chunk-size. Fig 4 shows that the average latency decreases as cache size increases, where average latency is 23 sec when no file has any content in the cache, and is 0 sec when the cache size is 4000 chunk-size since 4 chunks of each file can be in the cache. We note that the latency is a convex decreasing function of the cache size, depicting that our algorithm is effectively updating content in the cache and showing diminishing returns in a decrease of latency after reaching certain cache size.

**Evolution of Cache Content Placement:** We validate that our algorithm minimizes service latency by optimizing the cache content with respect to request arrival rates (the algorithm is minimizing latency bound  $\bar{U}_{i,t}$ , and  $\lambda_i$  and  $\Lambda_j$  is playing an important role in  $\bar{U}_{i,t}$ ), i.e., placing more file chunks that are frequently accessed (with a higher arrival rate) in the cache. We design a simulation which consists of 10 files, of size 100MB, and run our algorithm for three time bins, where each time bin has a different request arrival rate for the  $r = 10$  files as given in Table I. The arrows indicate the increase or decrease of the file arrival rate in the consecutive time bins. In this simulation, we plot the cache placement for the 10 files in steady state (after the algorithm converges) for each time bin. From Figure 5, we see that in the first time bin, as file 4 and file 9 had the highest request arrival rates, they had the highest numbers of chunks in the cache in the steady state. In the second time bin, as the arrival rates of files 4 and 9 decreased, and those of files 5 and 10 increased, now the files that had the same and higher request arrival rate were 1, 2, 6, and 7. Thus Fig 5 shows that in the second time bin these four files dominates the cache's steady state. In the third time bin, the files that had the highest arrival rate became files 2, 7, 4, and 9, and Fig. 5 also shows that cache is mainly holding the contents of file 2, 7, and 9. Due to different service rates of the servers and the randomized placement of contents in the servers, it is not always the case that files with the highest arrival rate need to be placed completely in the cache. Rather, it is important to determine how many chunks of each file should be placed in the cache.

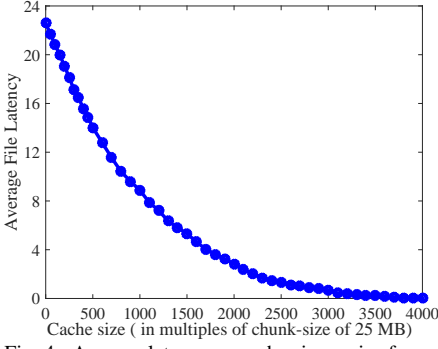


Fig. 4. Average latency as cache size varies from 0 to 4000 chunk-size, each data point represents average latency for  $r = 1000$  files in one time bin. This clearly demonstrates that the average latency decreases as cache size increases when all other parameters are fixed.

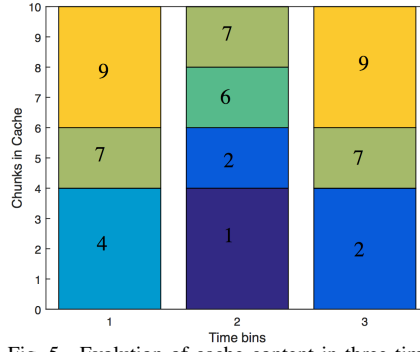


Fig. 5. Evolution of cache content in three time bins, each having different request arrival rates for the  $r = 10$  files. It shows that our algorithm is updating content held in the cache according to the intensity of workload of each file effectively.

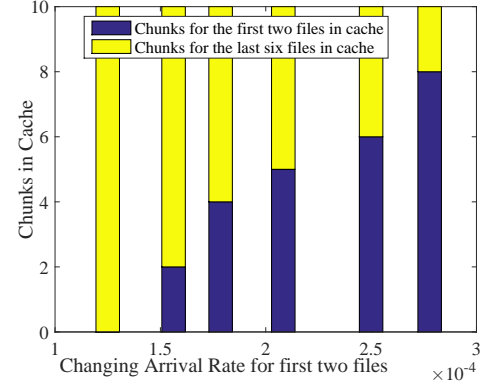


Fig. 6. The chunk placement depends not only on arrival rates but also on the placement of contents in the storage nodes and the service rates.

TABLE I  
TABLE OF ARRIVAL RATES FOR 10 FILES IN 3 TIME BINS.

Time Bin	File 1	File 2	File 3	File 4	File 5	File 6	File 7	File 8	File 9	File 10
1	0.000156	0.000156	0.000125	0.000167	0.000104	0.000156	0.000156	0.000125	0.000167	0.000104
2	0.000156	0.000156	0.000125	0.000125 ↓	0.000125 ↑	0.000156	0.000156	0.000125	0.000125 ↓	0.000125 ↑
3	0.000125 ↓	0.00025 ↑	0.000125	0.000167	0.000104	0.000125 ↓	0.00025 ↑	0.000125	0.000167	0.000104

**Impact of Content Placement and Arrival Rate:** We note that in addition to the arrival rates of the files, placement of content on the storage nodes influence the decision of cache placement. We consider 10 files, of size 100M, using (7,4) erasure code are placed on 12 servers as described in the simulation set up. The first three files are placed on the first seven servers while the rest of the files are placed on the last seven servers. Note that servers 6 and 7 host chunks for all files. For the arrival rates of the files, we fix the arrival rate for the last eight files such that the arrival rate of the third and fourth file is .0000962/sec, and of the last six files is 0.0001042/sec. The arrival rates for the first two files are assumed to be the same, and due to symmetry of files and placement, we consider four categories of content placement in the cache - contents for the first two files, contents for 3rd file, contents for 4th file, and the contents for the last six files. In the considered arrival rates in Figure 6, there was no chunk for the third and the fourth file placed in the cache due to low arrival rates. We note that we always assume the arrival rates of the first two files as the highest but since the servers on which they are placed have relatively lower average load, the arrival rate needs to be significantly higher for them to increase content in the cache. The six bars in Figure 6 correspond to arrival rates for the first two files of 0.0001250, 0.0001563, 0.0001786, 0.0002083, 0.0002500, and 0.0002778, respectively. At an arrival rate of .000125/sec for the first two files, there is no content for these files in the cache. Thus even though the arrival rate for the first two files is the largest, they do not need to be placed in the cache due to the fact that they are placed in the storage servers that are lightly loaded. As the arrival rate increases to .00015625/sec, two chunks corresponding to these files start to appear in the cache. Further increase of the arrival rate leads to more and more chunks in the cache. Thus, we note that the placement of chunks on the storage nodes, arrival rates, and service rates all play a role in allocating chunks of different files in the cache.

**Chunk Request Scheduling Evolution:** In order to see how the request scheduling evolves during each time bin, we

run the experiment with  $r = 1000$  objects, each of size 200 MB and using a (7,4) erasure code, and a total cache size of 62.5 GB. The average file request arrival rate for the two experiments is chosen as  $\lambda_i = 0.0225$  /sec and  $\lambda_i = 0.0384$ , respectively. We divide each time bin (100 sec) into 20 time slots, each with a length of 5 sec, and plot the number of chunks the client is getting from the cache and from the storage nodes in each time slot. Fig 7 shows that under both workloads, the number of chunks retrieved from the cache is smaller than that from OSDs. As the cache size is 62.5 GB, with chunk size 50 MB, it has a capacity of 1250 chunks, which means each file has more chunks in OSDs than that in cache on an average. Further, since the arrival rate of all file increases proportionally, the relative percentage of chunks retrieved from cache over 100s stays almost the same in the two cases, at about 33%.

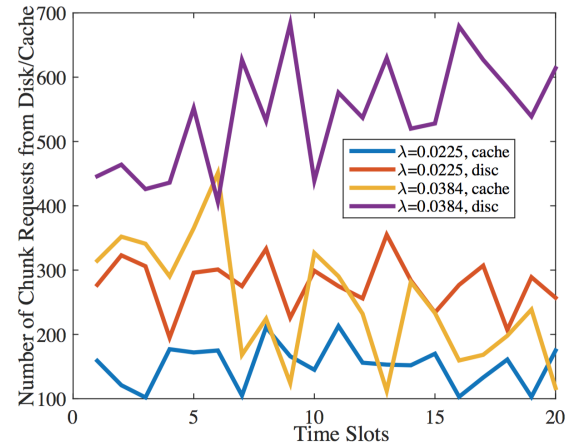


Fig. 7. Number of chunk requests that are sent to storage servers or cache with an average arrival rate  $\lambda_i = 0.0225$  /s and  $\lambda_i = 0.0384$ , respectively, where a time bin of length 100 s is divided into 20 equal time slots. The figure shows the dynamic nature of the number of chunks the client obtained from the cache/nodes.



### C. Ceph Testbed

We have validated the effectiveness of our latency analysis and the algorithm for joint optimization on average latency in erasure-coded storage systems with caching through simulation. In the rest of this section, we will evaluate the performance of our solution by implementing our algorithm in *Ceph* (the Jewel version with filestore), which is an open-source, object-based distributed storage system [25]. Details on Ceph storage API's and cache tier are given in Appendix A.

**Our Solution:** Our algorithm requires a caching tier that can store erasure-coded chunks in cache based on the proposed optimization. Since Ceph cache tier only provides replicated caching capability, we leverage it to implement an equivalent version of our functional caching policy. Typically we use a faster device (SSD) for cache, and by comparing the read latency from an OSD backed with HDD (shown in Table IV) and SSD (shown in Table V), we can see that service latency from cache can be negligible compared to that from back-end storage pool. As defined in Equation (1), the read latency of a file in erasure coded storage depends on the largest retrieval delay of the chunks (this is also true in Ceph). So the file latency depends only on the retrieval latency of chunks stored in back-end storage. Then the read process of file with an  $(n, k)$  code and  $d$  chunks in cache becomes equivalent to reading a file with  $(n, k - d)$  code (with the same chunk size) from back-end storage pool, while the latency of the  $d$  chunks in cache is ignored.

This allows us to implement our functional caching policy using different  $(n, k - d)$  codes and measure the resulting latency. More specifically, as the algorithm provides number of chunks of each file that are placed in the cache, we can dynamically adjust the equivalent erasure code in Ceph, by creating pools with the new equivalent erasure code and forward the file requests to the new pools. In this evaluation with Ceph, to be coherent with simulation setup, we create a Ceph cluster with 12 OSDs, using  $(7, 4)$  original erasure code with  $r = 1000$  files, which means as we put more chunks in cache, there will be 5 pools in total:  $(7, 4 - d)$  for  $d = 0, \dots, 4$ . And all of the five pools are backed by the same set of OSDs in the Ceph cluster. So that this environment provides 12 OSDs for object storage, each of the 1000 files has 7 chunks in the storage pool, and 0 to 4 chunks in the cache (changes with workload/timebin). So when a client tries to access a file, the file will be accessed from one of the five pools according to its equivalent code (calculated from the number of chunks in the cache of each file), from the algorithm at that time.

### D. Experiment Setup

**Testbed Configuration:** Throughout the evaluation, we compare the performance of our algorithm for optimal caching with a baseline: Ceph's caching tier, replicated caching with LRU. We have a test bed with 3 nodes for the Ceph object storage cluster, and another four nodes each with one virtual machine (VM) as 4 Ceph clients, each connected with 10G Ethernet network interface. All OSDs in this testbed is formatted with an ext4 filesystem. The setup for the two cases can be described as follows: (also shown in Fig 8)

1. **Optimal Erasure-Coded Caching:** Each node has 4 OSDs, and each OSD is backed with a HDD drive of size 900G. There are 5 pools created with the same set of 12 OSDs. Simulated cache capacity of 2500 chunks of 16 MB files.(1000 files has 4000 chunks in total, cache capacity 10G). Each node has one SSD drive that provides four 5G partitions for the journal of each OSD. The number of placement group for each pool is calculated from Equation (17), which is 256.

2. **Ceph Replicated Caching with LRU:** Cache tier consists of 2 OSDs, each backed with a 5G SSD partition, cache mode write-back with a capacity of 10G. Cache tier uses dual replication (close to the redundancy provided by  $(7, 4)$  erasure code). Storage tier has the same set as that in optimal erasure-coded caching. There is one pool created with  $(7, 4)$  erasure code on the 12 OSDs. The SSD partitions come from different SSD drives on one node. The number of placement group in cache tier is 128 from Equation (17). Journal settings and number of placement groups in the storage pool are the same with the optimal caching case.

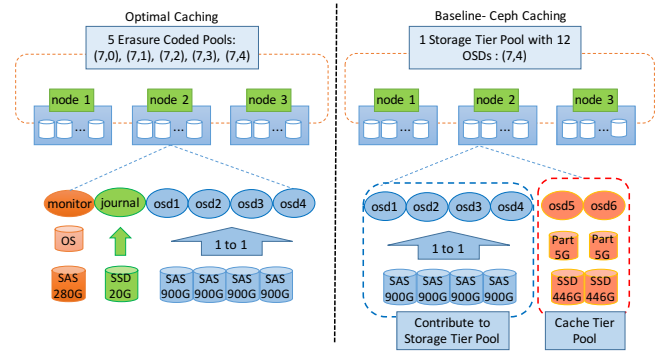


Fig. 8. Our Ceph Testbed with 3 nodes for proposed optimal caching and ceph's LRU caching as a baseline.

**Benchmarking Our Ceph Testbed:** We use COSBench 0.4.2 ([58], a benchmarking tool for cloud object storage developed by Intel) to benchmark the latency performance (response time in COSBench) in this evaluation. As file requests can be forward to any of the five storage pools according to the algorithm dynamically, to ensure that a file can be accessed from any pools at any time, we perform a write bench mark without cleaning-up to the pools before we run any read benchmarks. We modified COSBench to use existing pools instead of creating new pools (which is by default). We started 4 COSBench drives on the 4 VM clients. Each COSBench workload has 5 stages: Initial, Prepare, Main, Cleanup, and Dispose. For obtaining latency performance, the number of threads (workers in COSBench) in each driver in the first three stages is set to 1. The list of parameters is shown in Table II.

We use COSBench workload generator to emulate a real distributed storage workload, which is synthesized from actual production traces of a large-scale object storage system deployed in a commercial service provider. We studied the 24-hour object storage workload and choose the most popular object sizes (other similar object sizes are round up or down to the sizes in the table) for our test, and the average request arrival rate for each object in each size during the 24 hours is shown in Table III.

TABLE II  
TABLE OF COSBENCH DEFAULT CONFIGURATION

Number of Containers(pools)	5 or 1
Number of VM Clients	4
Object Size	4MB to 1GB
Access Pattern	Write, Read
Number of Workers	1
Rampup	30 s
Run time	1800 s

TABLE III  
TABLE OF 24-HOUR REAL STORAGE WORKLOAD.

Object Size	Average Request Arrival Rate
4MB	0.00029868
16MB	0.00010824
64MB	0.00051852
256MB	0.0000078
1GB	0.0000024

### E. Evaluation

Now we can evaluate the latency performance of our proposed optimal caching with different types of workloads, compared to Ceph's LRU caching as a baseline.

**Service Time Distribution:** In a Ceph erasure coded pool, a primary OSD receives all write operations, it encodes the object into  $n = k + m$  chunks and sends them to other OSDs (according to the value of  $n$ ). When an object is read from the erasure coded pool, it first checks the status of all OSDs that has a chunk of the required object, and then try to contact all such OSDs for retrieve a chunk. The decoding function can be called as soon as  $k$  chunks are read. While our service delay bound applies to arbitrary distribution and works for systems hosting any number of objects, we first run an experiment to understand actual service time distribution on our testbed, which is used as an input in the algorithm. To find the chunk service time distribution for different chunk sizes, we run a 100% write benchmark for all object sizes in Table III, into the (7,4) erasure coded pool in baseline testbed ( Ceph Replicated Caching with LRU) without cleaning-up for 900 seconds. Then, we run a 100% read benchmark for each object size with an average request arrival rate per object (among all objects of each size) as shown in III. For the five set experiments with five different object sizes, we run each set for 1800 seconds to get enough number of requests to take an average according to arrival rates in Table III. The chunk sizes are 1 MB, 4 MB, 16 MB, 64 MB, 256 MB, accordingly. We collect the service time of the chunks of different sizes at the OSDs and plot the results. Figure 9 depicts the Cumulative Distribution Function (CDF) of the chunk service times for read operations with different chunk sizes. We ignored the result for 1G object (256 MB chunk size) as in the figure as it is much larger than others and it would be hard to see the distribution of the first 4 if we add the result. The average service time at the OSD for read operations with chunk size 256 MB is 6758.06 milliseconds. Using the measured results, we get the mean and variance of service time of each chunk size as shown in Table IV. We use the obtained mean, and variance of the service time of chunks with different sizes to calculate moments of the service time distribution that are

TABLE IV  
MEAN AND VARIANCE OF SERVICE TIME FOR DIFFERENT CHUNK SIZES IN MILLISECOND.

Chunk Size	Mean Service Time	Variance of Service Time
1MB	6.6696	0.0963
4MB	35.8800	2.6925
16MB	147.8462	388.9872
64MB	355.0800	1256.6100
256MB	6758.06	554180

used to find cache placement in our algorithm.

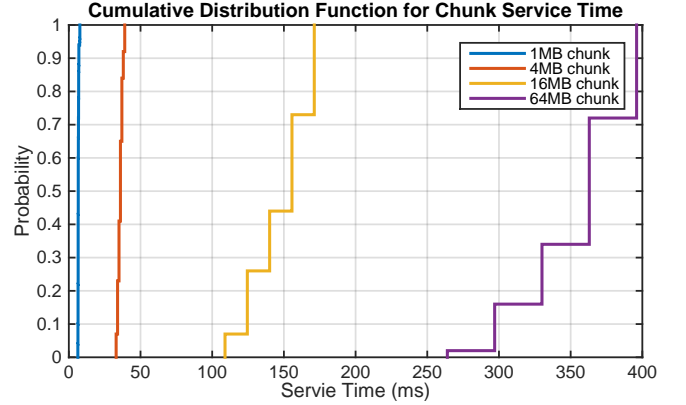


Fig. 9. Actual service latency distribution for different chunk sizes from the Ceph testbed with 12 OSDs without cache tier.

**Evaluation of the Performance of Our Solution:** As we have validated the algorithm in aspects of convergence and evolution of dynamic caching, we now evaluate the effectiveness of the proposed algorithm (in terms of overall latency improvement) in real storage systems with the real workload. For experiment in optimal caching, as we introduced earlier, we have five pools set up with erasure code (7,0), (7,1), (7,2), (7,3), (7,4), where pool (7,0) means all the four required chunks are stored in cache, which will typically be some local SAS or NVMe SSDs in the case of optimal caching. We run a read test to the SSD drives for all object sizes in Table III and get the average latency numbers for each object size shown in Table V. As compared with average service time of the same chunk size from an OSD (HDD as backend) shown in Table IV, we can see that the latency of a chunk from cache can be negligible when the read latency of the whole object is measured. This is because the latter depends on the chunk read latency from an OSD in the Ceph cluster, which motivates the use of equivalent codes in the caching schemes.

TABLE V  
READ LATENCY OF DIFFERENT CHUNK SIZES FROM CACHE (SAS SSD) IN OPTIMAL CACHING (MS).

Chunk Size	Read Latency From Cache
1 MB	1.86619
4 MB	7.35639
16 MB	30.4927
64 MB	97.0968
256 MB	349.133

In each time bin, given the workload and the total number of active objects (for simplicity of the algorithm, we set this number to 1000, which is also very close to the number of

objects that have been accessed during workload run time for 30-min in the real workload) in the Ceph object storage cluster, the optimization algorithm provides which object belongs to which pool (an object-pool map), and the read request arrival rate of each object in each pool. In order to evaluate the effectiveness of the optimal caching algorithm with various object sizes, first, we run a 100% write benchmark into the erasure coded pools according to the object-pool map from the optimization for each object size in Table III. Then, for each object size, we run read benchmarks to the erasure-code pools (according to the object-pool map from the optimization) for 1800 seconds. The average read request arrival rate per object in each pool can be calculated as the number of objects in the pool from the algorithm times the average read request arrival rate for each object size in Table III. After the experiment is done, we get the average read latency for objects in each erasure coded pool, and then calculate the average latency over the 1800 seconds time bin as,

$$\sum_i \frac{\text{number of objects in pool } i * \text{average latency in pool } i}{\text{total number of objects in all pools}}. \quad (16)$$

For evaluation of Ceph's LRU replicated caching as a baseline, first, we run 100% write benchmark (with the same number of objects we used in optimal caching) for each object size in Table III to the (7,4) erasure-code storage pool with a replicated cache tier overlay. Then for each object size, we perform read workloads to this pool with an average read request arrival rate per object as shown in Table III for 1800 seconds. Aggregated request arrival rate in this test for the (7,4) pool would be the request arrival rate in Table III times the total number of active objects (1000 in this case). In both the cases, cache capacity is fixed and set to be 10GB. This cache can hold 10000 chunks of 4 MB objects, 2500 chunks of 16 MB objects, 625 chunks of 64 MB objects, 156 chunks of 256 MB objects, and 39 chunks of 1 GB objects when using (7,4) erasure code. For example, each object of size 16 MB can have  $2500/1000=2.5$  chunks on average stored in the cache in both optimal caching and Ceph cache-tier. However, as each object of each size has a different read request arrival rate (value shown in III is an average of all files of the same size), placing two or three chunks in the cache would help reduce latency significantly, thus leading to uneven placement of chunks of the objects in the cache. For this experiment with baseline, we get the average read latency for objects in (7,4) erasure coded pool. We compare the average latency to retrieve an object in optimal caching and Ceph's LRU caching. Fig 10 shows that average latency increases as the object size increases, which we can see from our latency bound since the number of chunks that the cache can hold decreases as object size increases. The figure also shows that our caching improves latency as compared to Ceph's LRU caching as a baseline, which is using dual replication in its cache tier, by 26% on average. This improvement becomes more and more significant as the number of objects increase, which shows that our dynamic optimal chunk placement in the cache is more effective than traditional LRU cache. The improvement with increased file size is because the load on the system increases thus obtaining more latency advantages

with caching. For a lightly-loaded system, the latency is small and the improvement in caching may not be that significant.

Fig. 10 shows that our algorithm with optimal caching significantly improves latency with a reasonable number of objects and cache sizes. Fig. 10 also compares the analytically optimized expression of latency from the simulations and the experimentally observed latency where we see that the analytical result is an upper bound and matches well with the experimental results.

Next, we evaluate the performance of optimal caching with various workload intensities. We fix the object size at 64 MB (we still use 1000 objects) and vary the read request arrival rate for these objects and evaluate the latency performance. Queuing latency would not be a dominant issue if the workload is not intense. Thus, we fully load the Ceph cluster with much higher request arrival rates as compared to those in Table III. The aggregate read request arrival rates (Single object request arrival rate times the number of objects) used in this test are 0.5, 1.0, 2.0, 4.0, and 8.0. In the optimal caching case, for each workload (request arrival rate), we perform the read benchmark to the five erasure coded pools according to the object-pool map from the optimization with this request arrival rate for 1800 seconds. For Ceph's LRU caching, we perform read benchmarks to the (7,4) erasure coded pool with the same set of arrival rates in the case of optimal caching for 1800 seconds, and obtain the average access latency per object. The cache size for both cases is still fixed at 10 GB. Actual average service latency of objects for each workload is shown by a bar plot in Figure 11. In this experiment we also compare the results for the optimal caching scheme and Ceph's LRU caching as a baseline. Fig 11 shows that our optimal caching algorithm outperforms Ceph's LRU caching in terms of average latency for all workload intensities in the setup. The proposed algorithm gives an average 23.86% reduction in latency. Thus, our algorithm with optimal caching can mitigate traffic contention and reduce latency very efficiently compared to Ceph's LRU caching, this also provides a guideline for designing caching schemes under very heavy workloads.

## VI. CONCLUSIONS

In this paper, we propose functional caching, a novel approach to create erasure-coded cache chunks that maintain MDS code property along with existing data chunks. It outperforms exact caching schemes and provides a higher degree of freedom in file access and request scheduling. We quantify an upper bound on the mean service latency in closed-form for erasure-coded storage systems with functional caching, for arbitrary chunk placement and service time distributions. A cache optimization problem is formulated and solved using an efficient heuristic algorithm. Numerical results and prototype in an open-source cloud storage validate significant service latency reduction using functional caching.

This paper assumes that a rate monitoring/prediction mechanism is available to detect the rate changes. This can be an online predictive model [59] or a simple sliding-window-based method, which continuously measures the average request arrival, and when the difference of request arrival rate of an object between two different time windows is greater than a threshold, a new time bin is considered to trigger a new



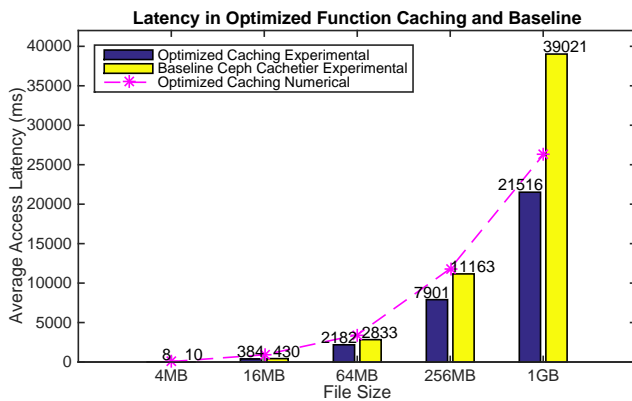


Fig. 10. Average latency in optimal erasure coded caching for different object sizes during 1800 seconds benchmark run time, compared with Ceph's LRU replicated caching as a baseline. Cache capacity fixed at 10 GB in both cases. Figure shows significant latency improvement in optimal caching.

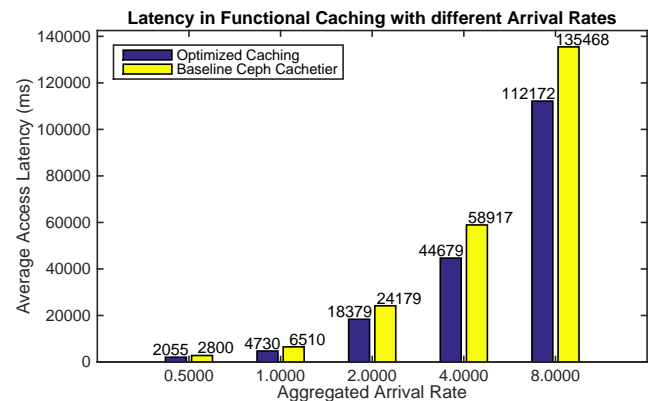


Fig. 11. Comparison of average latency in optimal erasure coded caching and Ceph's LRU replicated caching for different intensity of read workload during 1800 seconds benchmark run time. Cache capacity fixed at 10 GB in both cases. Figure shows significant latency improvement in optimal caching.

round of optimization. Finding a robust algorithm that can automatically adjust to such changes is an open problem.

## REFERENCES

- [1] V. Aggarwal, Y. F. R. Chen, T. Lan, and Y. Xiang, "Sprout: A functional caching approach to minimize service latency in erasure-coded storage," in *2016 IEEE 36th International Conference on Distributed Computing Systems (ICDCS)*, June 2016, pp. 753–754.
- [2] H. Weatherspoon and J. Kubiatowicz, "Erasure coding vs. replication: A quantitative comparison," in *Revised Papers from the First International Workshop on Peer-to-Peer Systems*, ser. IPTPS '01. Springer-Verlag, 2002.
- [3] A. Dimakis, P. Godfrey, Y. Wu, M. Wainwright, and K. Ramchandran, "Network coding for distributed storage systems," *Information Theory, IEEE Transactions on*, vol. 56, no. 9, pp. 4539–4551, Sept 2010.
- [4] M. Sathiamoorthy, M. Asteris, D. Papailiopoulos, A. G. Dimakis, R. Vadali, S. Chen, and D. Borthakur, "Xoring elephants: Novel erasure codes for big data," in *Proceedings of the 39th international conference on Very Large Data Bases.*, 2013.
- [5] C. Huang, H. Simitci, Y. Xu, A. Ogus, B. Calder, P. Gopalan, J. Li, and S. Yekhanin, "Erasure coding in windows azure storage," in *Proceedings of the 2012 USENIX Conference on Annual Technical Conference*, ser. USENIX ATC'12. USENIX Association, 2012.
- [6] A. Fikes, "Storage architecture and challenges (talk at the google faculty summit)," <http://bit.ly/nUylRW>, Tech. Rep., 2010.
- [7] L. Huang, S. Pawar, H. Zhang, and K. Ramchandran, "Codes can reduce queueing delay in data centers," in *Information Theory Proceedings (ISIT)*, 2012 IEEE International Symposium on, July 2012, pp. 2766–2770.
- [8] G. Joshi, Y. Liu, and E. Soljanin, "On the delay-storage trade-off in content download from coded distributed storage systems," *Selected Areas in Communications, IEEE Journal on*, vol. 32, no. 5, pp. 989–997, May 2014.
- [9] S. Chen, Y. Sun, U. Kozat, L. Huang, P. Sinha, G. Liang, X. Liu, and N. Shroff, "When queuing meets coding: Optimal-latency data retrieving scheme in storage clouds," in *Proceedings of IEEE Infocom*, 2014.
- [10] N. Shah, K. Lee, and K. Ramchandran, "The mds queue: analyzing latency performance of codes and redundant requests," *arXiv:1211.5405*, Nov 2012.
- [11] Y. Xiang, T. Lan, V. Aggarwal, and Y. F. R. Chen, "Joint latency and cost optimization for erasure-coded data center storage," *SIGMETRICS Perform. Eval. Rev.*, vol. 42, no. 2, pp. 3–14, Sep. 2014. [Online]. Available: <http://doi.acm.org/10.1145/2667522.2667524>
- [12] —, "Joint latency and cost optimization for erasure-coded data center storage," *IEEE/ACM Transactions on Networking*, vol. 24, no. 4, pp. 2443–2457, Aug 2016.
- [13] Y. Xiang, T. Lan, V. Aggarwal, and Y.-F. Chen, "Multi-tenant latency optimization in erasure-coded storage with differentiated services," in *Distributed Computing Systems (ICDCS), 2015 IEEE 35th International Conference on*, June 2015, pp. 790–791.
- [14] Y. Xiang, V. Aggarwal, Y.-F. Chen, and T. Lan, "Taming latency in data center networking with erasure coded files," in *Cluster, Cloud and Grid Computing (CCGrid), 2015 15th IEEE/ACM International Symposium on*, May 2015, pp. 241–250.
- [15] R. Pedarsani, M. A. Maddah-Ali, and U. Niesen, "Online coded caching," in *IEEE International Conference on Communications, ICC 2014, Sydney, Australia, June 10-14, 2014*, 2014, pp. 1878–1883.
- [16] S. J. Nadgowda, R. C. Sreenivas, S. Gupta, N. Gupta, and A. Verma, *Service-Oriented Computing: 12th International Conference, ICSOC 2014, Paris, France, November 3-6, 2014. Proceedings.* Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, ch. C2P: Co-operative Caching in Distributed Storage Systems, pp. 214–229.
- [17] F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber, "Bigtable: A distributed storage system for structured data," *ACM Trans. Comput. Syst.*, vol. 26, no. 2, pp. 4:1–4:26, Jun. 2008.
- [18] Q. Zhu, A. Shankar, and Y. Zhou, "Pb-lru: A self-tuning power aware storage cache replacement algorithm for conserving disk energy," in *Proceedings of the 18th Annual International Conference on Supercomputing*, ser. ICS '04. New York, NY, USA: ACM, 2004, pp. 79–88.
- [19] M. Naik, F. Schmuck, and R. Tewari, "Read and write requests to partially cached files," Aug. 4 2015, uS Patent 9,098,413. [Online]. Available: <http://www.google.com/patents/US9098413>
- [20] H. B. T. K. A. A. R. Eno Thereska, Greg O'Shea, "Software-defined caching: Managing caches in multi-tenant data centers," in *ACM Symposium on Cloud Computing (SOCC) 2015*. ACM Association for Computing Machinery, August 2015. [Online]. Available: <https://www.microsoft.com/en-us/research/publication/software-defined-caching-managing-caches-in-multi-tenant-data-centers/>
- [21] Y.-J. Hong and M. Thottethodi, "Understanding and mitigating the impact of load imbalance in the memory caching tier," in *Proceedings of the 4th Annual Symposium on Cloud Computing*, ser. SOCC '13. New York, NY, USA: ACM, 2013, pp. 13:1–13:17. [Online]. Available: <http://doi.acm.org/10.1145/2523616.2525970>
- [22] Y. Cheng, A. Gupta, and A. R. Butt, "An in-memory object caching framework with adaptive load balancing," in *Proceedings of the Tenth European Conference on Computer Systems*, ser. EuroSys '15. New York, NY, USA: ACM, 2015, pp. 4:1–4:16. [Online]. Available: <http://doi.acm.org/10.1145/2741948.2741967>
- [23] L. Wang, S. Bayhan, and J. Kangasharju, "Optimal chunking and partial caching in information-centric networks," *Comput. Commun.*, vol. 61, no. C, pp. 48–57, May 2015. [Online]. Available: <http://dx.doi.org/10.1016/j.comcom.2014.12.009>
- [24] D. Bertsimas, K. Natarajan, and C.-P. Teo, "Tight bounds on expected order statistics," *Probab. Eng. Inf. Sci.*, vol. 20, no. 4, pp. 667–686, Oct. 2006.
- [25] S. A. Weil, S. A. Brandt, E. L. Miller, D. D. E. Long, and C. Maltzahn, "Ceph: A scalable, high-performance distributed file system," in *Proceedings of the 7th Symposium on Operating Systems Design and Implementation*, ser. OSDI '06. Berkeley, CA, USA: USENIX Association, 2006, pp. 307–320. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1298455.1298485>




- [26] A. Kumar, R. Tandon, and T. Clancy, "On the latency and energy efficiency of distributed storage systems," *IEEE Transactions on Cloud Computing*, vol. PP, no. 99, pp. 1–1, 2015.
- [27] M. Noori, E. Soljanin, and M. Ardakani, "On storage allocation for maximum service rate in distributed storage systems," in *2016 IEEE International Symposium on Information Theory (ISIT)*, July 2016, pp. 240–244.
- [28] G. Joshi, E. Soljanin, and G. Wornell, "Queues with redundancy: Latency-cost analysis," *SIGMETRICS Perform. Eval. Rev.*, vol. 43, no. 2, pp. 54–56, Sep. 2015. [Online]. Available: <http://doi.acm.org/10.1145/2825236.2825258>
- [29] V. Aggarwal and T. Lan, "Tail index for a distributed storage system with pareto file size distribution," *CoRR*, vol. abs/1607.06044, 2016. [Online]. Available: <http://arxiv.org/abs/1607.06044>
- [30] F. Meyer auf der Heide, B. Vöcking, and M. Westermann, "Caching in networks (extended abstract)," in *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms, January 9-11, 2000, San Francisco, CA, USA., 2000*, pp. 430–439.
- [31] M. R. Korupolu, C. G. Plaxton, and R. Rajaraman, "Placement algorithms for hierarchical cooperative caching," in *Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA '99, 1999, pp. 586–595.
- [32] S. Jiang and X. Zhang, "Making lru friendly to weak locality workloads: a novel replacement algorithm to improve buffer cache performance," *Computers, IEEE Transactions on*, vol. 54, no. 8, pp. 939–952, Aug 2005.
- [33] N. Megiddo and D. S. Modha, "Arc: A self-tuning, low overhead replacement cache," in *Proceedings of the 2Nd USENIX Conference on File and Storage Technologies*, ser. FAST '03, 2003, pp. 115–130.
- [34] Y. Zhou, J. Philbin, and K. Li, "The multi-queue replacement algorithm for second level buffer caches," in *Proceedings of the General Track: 2001 USENIX Annual Technical Conference*, 2001, pp. 91–104.
- [35] S. Liang, K. Chen, S. Jiang, and X. Zhang, "Cost-aware caching algorithms for distributed storage servers," in *Distributed Computing*, ser. Lecture Notes in Computer Science, A. Pelc, Ed. Springer Berlin Heidelberg, 2007, vol. 4731, pp. 373–387.
- [36] S. Angelopoulos, R. Dorriv, and A. López-Ortiz, "On the separation and equivalence of paging strategies," in *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA '07, 2007, pp. 229–237.
- [37] D. S. M. Elisha J. Rosensweig and J. Kurose, "On the steady-state of cache networks," in *Proceedings of the IEEE INFOCOM*, 2013.
- [38] J. Zhang, X. Lin, and X. Wang, "Coded caching under arbitrary popularity distributions," in *2015 Information Theory and Applications Workshop, ITA 2015, San Diego, CA, USA, February 1-6, 2015*, 2015, pp. 98–107.
- [39] K. Rashmi, M. Chowdhury, J. Kosaian, I. Stoica, and K. Ramchandran, "Ec-cache: load-balanced, low-latency cluster caching with online erasure coding," in *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*. USENIX Association, 2016, pp. 401–417.
- [40] Q. Huang, K. Birman, R. van Renesse, W. Lloyd, S. Kumar, and H. C. Li, "An analysis of facebook photo caching," in *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*, ser. SOSP '13, 2013, pp. 167–181.
- [41] N. Bronson, Z. Amsden, G. Cabrera, P. Chakka, P. Dimov, H. Ding, J. Ferris, A. Giardullo, S. Kulkarni, H. Li, M. Marchukov, D. Petrov, L. Puzar, Y. J. Song, and V. Venkataramani, "Tao: Facebook's distributed data store for the social graph," in *Presented as part of the 2013 USENIX Annual Technical Conference (USENIX ATC 13)*, San Jose, CA, 2013, pp. 49–60.
- [42] B. Fitzpatrick, "Distributed caching with memcached," *Linux J.*, vol. 2004, no. 124, pp. 5–, Aug. 2004.
- [43] G. Chockler, G. Laden, and Y. Vigfusson, "Data caching as a cloud service," in *Proceedings of the 4th International Workshop on Large Scale Distributed Systems and Middleware*, ser. LADIS '10, 2010, pp. 18–21.
- [44] —, "Design and implementation of caching services in the cloud," *IBM Journal of Research and Development*, vol. 55, no. 6, pp. 9:1–9:11, Nov 2011.
- [45] U. Niesen and M. A. Maddah-Ali, "Coded caching for delay-sensitive content," *CoRR*, vol. abs/1407.4489, 2014.
- [46] N. Karamchandani, U. Niesen, M. Maddah-Ali, and S. Diggavi, "Hierarchical coded caching," in *Information Theory (ISIT), 2014 IEEE International Symposium on*, June 2014, pp. 2142–2146.
- [47] M. Maddah-Ali and U. Niesen, "Fundamental limits of caching," *Information Theory, IEEE Transactions on*, vol. 60, no. 5, pp. 2856–2867, May 2014.
- [48] U. Niesen and M. Maddah-Ali, "Coded caching with nonuniform demands," in *Computer Communications Workshops (INFOCOM WK-SHPS), 2014 IEEE Conference on*, April 2014, pp. 221–226.
- [49] S. P. Shariatpanahi, A. S. Motahari, and B. H. Khalaj, "Multi-server coded caching," *CoRR*, vol. abs/1503.00265, 2015.
- [50] A. Dimakis, V. Prabhakaran, and K. Ramchandran, "Distributed data storage in sensor networks using decentralized erasure codes," in *Signals, Systems and Computers, 2004. Conference Record of the Thirty-Eighth Asilomar Conference on*, vol. 2, Nov 2004, pp. 1387–1391 Vol.2.
- [51] M. Aguilera, R. Janakiraman, and L. Xu, "Using erasure codes efficiently for storage in a distributed system," in *Dependable Systems and Networks, 2005. DSN 2005. Proceedings. International Conference on*, June 2005, pp. 336–345.
- [52] Q. Lv, P. Cao, E. Cohen, K. Li, and S. Shenker, "Search and replication in unstructured peer-to-peer networks," in *Proceedings of the 16th International Conference on Supercomputing*, ser. ICS '02. New York, NY, USA: ACM, 2002, pp. 84–95. [Online]. Available: <http://doi.acm.org/10.1145/514191.514206>
- [53] W. W. Peterson and E. Weldon, "Error-correcting codes," in *the MIT Press*, March 1972.
- [54] A. Abdelkefi and Y. Jiang, "A structural analysis of network delay," in *Communication Networks and Services Research Conference (CNSR), 2011 Ninth Annual*, May 2011, pp. 41–48.
- [55] H. Weatherspoon and J. Kubiatowicz, "Erasure coding vs. replication: A quantitative comparison," in *Peer-to-Peer Systems*, ser. Lecture Notes in Computer Science, P. Druschel, F. Kaashoek, and A. Rowstron, Eds. Springer Berlin Heidelberg, 2002, vol. 2429, pp. 328–337.
- [56] A. Downey, "The structural cause of file size distributions," in *Modeling, Analysis and Simulation of Computer and Telecommunication Systems, 2001. Proceedings. Ninth International Symposium on*, 2001, pp. 361–370.
- [57] F. Paganini, A. Tang, A. Ferragut, and L. Andrew, "Network stability under alpha fair bandwidth allocation with general file size distribution," *Automatic Control, IEEE Transactions on*, vol. 57, no. 3, pp. 579–591, March 2012.
- [58] Q. Zheng, H. Chen, Y. Wang, J. Duan, and Z. Huang, "Cosbench: A benchmark tool for cloud object storage services," in *Proceedings of CLOUD 2012*, ser. CLOUD '12. Washington, DC, USA: IEEE Computer Society, 2012, pp. 998–999. [Online]. Available: <http://dx.doi.org/10.1109/CLOUD.2012.52>
- [59] D. Shen and J. L. Hellerstein, "Predictive models for proactive network management: Application to a production web server," in *Network Operations and Management Symposium, 2000. NOMS 2000. 2000 IEEE/IFIP*. IEEE, 2000, pp. 833–846.
- [60] J. S. Plank, S. Simmerman, and C. D. Schuman, "Jerasure: A library in c/c++ facilitating erasure coding for storage applications," Tech. Rep., 2007.
- [61] T. Angell, "The farkas-minkowski theorem," [www.math.udel.edu/~angell/Opt/farkas.pdf](http://www.math.udel.edu/~angell/Opt/farkas.pdf), Tech. Rep., 2002.



**Vaneet Aggarwal (S'08 - M'11 - SM15)** received the B.Tech. degree in 2005 from the Indian Institute of Technology, Kanpur, India, and the M.A. and Ph.D. degrees in 2007 and 2010, respectively from Princeton University, Princeton, NJ, USA, all in Electrical Engineering.

He is currently an Assistant Professor at Purdue University, West Lafayette, IN. Prior to this, he was a Senior Member of Technical Staff Research at AT&T Labs-Research, NJ, and an Adjunct Assistant Professor at Columbia University, NY. His research interests are in applications of statistical, algebraic, and optimization techniques to distributed storage systems, machine learning, and wireless systems. Dr. Aggarwal was the recipient of Princeton University's Porter Ogden Jacobus Honorary Fellowship in 2009. In addition, he received AT&T Key Contributor award in 2013, AT&T Vice President Excellence Award in 2012, and AT&T Senior Vice President Excellence Award in 2014.




Scientist and a Vice Chair of the International World Wide Web Conferences Steering Committee (IW3C2). He also serves on the editorial board of IEEE Internet Computing.

**Yih-Farn (Robin) Chen** is a Director of Inventive Science, leading the Cloud Platform Software Research department at AT&T Labs - Research. His current research interests include cloud computing, software-defined storage, mobile computing, distributed systems, World Wide Web, and IPTV. He holds a Ph.D. in Computer Science from University of California at Berkeley, an M.S. in Computer Science from University of Wisconsin, Madison, and a B.S. in Electrical Engineering from National Taiwan University. Robin is an ACM Distinguished



2009 IEEE GLOBECOM Best Paper Award, and the 2012 INFOCOM Best Paper Award.

**Tian Lan (S'03-M'10)** received the B.A.Sc. degree from the Tsinghua University, China in 2003, the M.A.Sc. degree from the University of Toronto, Canada, in 2005, and the Ph.D. degree from the Princeton University in 2010. Dr. Lan is currently an Associate Professor of Electrical and Computer Engineering at the George Washington University. His research interests include cloud resource optimization, mobile networking, storage systems and cyber security. Dr. Lan received the 2008 IEEE Signal Processing Society Best Paper Award, the



**Yu Xiang** received the B.A.Sc. degree from Harbin Institute of Technology in 2010, and the Ph.D. degree from George Washington University in 2015, both in Electrical Engineering. She is now a senior inventive scientist at AT&T Labs-Research. Her current research interest are in cloud resource optimization, distributed storage systems and cloud storage charge-back.

## APPENDIX CEPH TESTBED

The underlying foundation of Ceph is the Reliable Autonomic Distributed Object Store (RADOS). Ceph provides three types of storage API's on top of RADOS: (i) Object Storage: the fundamental RADOS object-based storage system. (ii) Block Storage: RADOS Block Device (RBD) that stores block device images as objects. (iii) Filesystem: this runs on top of the object storage with a metadata server that maps the directories and file names of the file system to objects stored within RADOS clusters. We use Ceph object storage in our testbed for our joint optimization algorithm with caching in erasure coded storage because i) Object storage is a good fit for large-scale cloud storage systems since it stores data more efficiently. ii) In terms of resilience, Ceph object storage provides both replicated and erasure-coded storage, while erasure-coded block storage on top of RADOS is not supported in Ceph yet.

**Ceph Object Storage Cluster:** A Ceph object storage cluster consists of two types of daemons: (i) Ceph monitor, which keeps a master copy of the cluster map and monitors the status of daemons in the cluster. (ii) Ceph object storage device(OSD), which is a physical or logical storage device used to store data as objects. Each OSD has two components: journal and data, during a write operation, it writes to the journal first and then writes data from journal to OSD data drive. Ceph stores data within pools, a pool consists of a bunch of OSDs, and Ceph users can specify which bunch of OSDs forms a certain pool through CRUSH rules. CRUSH is an algorithm Ceph used to pseudo-randomly store and retrieve data in OSDs with a uniform distribution of data across the cluster. CRUSH has a cluster map which contains a list of OSDs, and each OSD's physical location, and a list of rules that tell CRUSH how it should encode/replicate data in pools. To reduce the amount of metadata when storing data in OSDs, Ceph maps objects to placement groups, which are fragments of an object pool that place objects as a group into OSDs. The number of placement groups in an  $(n=k+m, k)$  erasure coded pool can be calculated as:

$$\text{Number of Placement Groups} = \frac{\text{Number of OSDs} * 100}{m} \quad (17)$$

**Erasure-coded Storage in Ceph:** When creating an erasure coded pool in Ceph, users need to specify erasure code parameters, the number of placement groups, OSDs belong to this pool, and failure domain of the pool (OSD, or host) etc., in an erasure-code profile. Ceph object storage is using *jerasure plugin* (a generic and flexible plugin that encapsulates the Jerasure code library) for default erasure code set up [60]. For an erasure-code profile with *jerasure plugin*, we specify erasure code with  $k$ , which is the number of data chunks, and  $m$ , which is the number of coded chunks, thus forming an  $(k + m, k)$  erasure code.

**Ceph Cache Tier:** Ceph provides a cache tiering service to improve IO performance for storage devices. Users can create a Ceph cache tier as an overlay to the existing storage pool. Usually, the cache tier is a pool with faster but expensive storage devices (mostly SAS SSDs or Nvme SSDs). The

existing back-end storage pool could be either an erasure-coded or replicated pool composed of slower and cheaper storage devices. However, the cache tier pool has to be replicated as Ceph caching does not support erasure-coded storage yet. For a pool with cache tier, all IO traffic will be routed to the cache tier first. In the case of a cache miss, the data will be promoted from the storage tier to the cache tier. And the tiering agent will evict the least recently used (LRU) data from cache tier and flushes them to storage tier.

#### APPENDIX PROOF OF LEMMA 1

The key idea of probabilistic scheduling is to get the file  $i$  from set  $\mathcal{A}_i$ , which is chosen probabilistically, and thus is given by the parameters  $\mathbb{P}(\mathcal{A}_i) \forall \mathcal{A}_i \subseteq \mathcal{S}_i$  and  $\forall i$ , which involves  $\sum_i \binom{n_i}{k_i}$  decision variables. While it appears prohibitive computationally, the authors of [12] demonstrated that the optimization can be transformed into an equivalent form, which only requires  $\sum_i n_i$  variables. The key idea is to show that it is sufficient to consider the conditional probability (denoted by  $\pi_{i,j}$ ) of selecting a node  $j$ , given that a batch of  $k_i$  chunk requests of file  $i$  are dispatched. It is easy to see that for given  $\mathbb{P}(\mathcal{A}_i)$ , we can derive  $\pi_{i,j}$  by

$$\pi_{i,j} = \sum_{\mathcal{A}_i: \mathcal{A}_i \subseteq \mathcal{S}_i} \mathbb{P}(\mathcal{A}_i) \cdot \mathbf{1}_{\{j \in \mathcal{A}_i\}}, \quad \forall i \quad (18)$$

where  $\mathbf{1}_{\{j \in \mathcal{A}_i\}}$  is an indicator function which equals to 1 if node  $j$  is selected by  $\mathcal{A}_i$  and 0 otherwise. The equivalence proof of the probability over sets  $\mathcal{A}_i$  and the nodes  $j$  follows using Farkas-Minkowski Theorem [61].

Using this result, it is sufficient to study probabilistic scheduling via conditional probabilities  $\pi_{i,j}$ , which greatly simplifies the analysis. In particular, it is easy to verify that under our model, the arrival of chunk requests at node  $j$  form a Poisson Process with rate  $\Lambda_j = \sum_i \lambda_i \pi_{i,j}$ , which is the superposition of  $r$  Poisson processes each with rate  $\lambda_i \pi_{i,j}$ ,  $\mu_j$  is the service rate of node  $j$ . The resulting queuing system under probabilistic scheduling is stable if all local queues are stable.

Let  $\mathbf{Q}_{max}$  be the maximum of waiting time  $\{\mathbf{Q}_j, j \in \mathcal{A}_i\}$ . We first show that  $\mathbf{Q}_{max}$  is upper bounded by the following inequality for arbitrary  $z \in \mathbb{R}$ :

$$\mathbf{Q}_{max} \leq z + [\mathbf{Q}_{max} - z]^+ \leq z + \sum_{j \in \mathcal{A}_i} [\mathbf{Q}_j - z]^+, \quad (19)$$

where  $[a]^+ = \max\{a, 0\}$  is a truncate function. Now, taking the expectation on both sides of (19), we have

$$\begin{aligned} \mathbb{E}[\mathbf{Q}_{max}] &\leq z + \mathbb{E} \left[ \sum_{j \in \mathcal{A}_i} [\mathbf{Q}_j - z]^+ \right] \\ &= z + \mathbb{E} \left[ \sum_{j \in \mathcal{A}_i} \frac{1}{2} (\mathbf{Q}_j - z + |\mathbf{Q}_j - z|) \right] \\ &= z + \mathbb{E}_{\mathcal{A}_i} \left[ \sum_{j \in \mathcal{A}_i} \frac{1}{2} (\mathbb{E}[\mathbf{Q}_j] - z + \mathbb{E}|\mathbf{Q}_j - z|) \right], \\ &= z + \sum_{j \in \mathcal{A}_i} \frac{\pi_{i,j}}{2} (\mathbb{E}[\mathbf{Q}_j] - z + \mathbb{E}|\mathbf{Q}_j - z|), \quad (20) \end{aligned}$$

where  $\mathbb{E}_{\mathcal{A}_i}$  denotes the expectation over randomly selected  $k_i$  storage nodes in  $\mathcal{A}_i \subseteq \mathcal{S}$  according to probabilities  $\pi_{i,1}, \dots, \pi_{i,m}$ . From Cauchy-Schwarz inequality, we have

$$\mathbb{E}|\mathbf{Q}_j - z| \leq \sqrt{(\mathbb{E}[\mathbf{Q}_j] - z)^2 + \text{Var}[\mathbf{Q}_j]}. \quad (21)$$

Combining (20) and (21), we obtain

$$\begin{aligned} \bar{T}_i &\leq \min_{z \in \mathbb{R}} \left\{ z + \sum_{j \in \mathcal{S}_i} \frac{\pi_{i,j}}{2} (\mathbb{E}[\mathbf{Q}_j] - z) \right. \\ &\quad \left. + \sum_{j \in \mathcal{S}_i} \frac{\pi_{i,j}}{2} \left[ \sqrt{(\mathbb{E}[\mathbf{Q}_j] - z)^2 + \text{Var}[\mathbf{Q}_j]} \right] \right\}. \quad (22) \end{aligned}$$

Using Pollaczek-Khinchin transform [7] to obtain the mean and variance of M/G/1 queueing delays, we obtain the result as in the statement of the Lemma.