

Coding for Fast Content Download

Gauri Joshi
EECS Dept., MIT
Cambridge, MA 02139, USA
Email: gauri@mit.edu

Yanpei Liu
ECE Dept., Univ. Wisconsin Madison
Madison, WI, 53705, USA
Email: yliu73@wisc.edu

Emina Soljanin
Bell Labs, Alcatel-Lucent
Murray Hill NJ 07974, USA
Email: emina@bell-labs.com

Abstract— We study the fundamental trade-off between storage and content download time. We show that the download time can be significantly reduced by dividing the content into chunks, encoding it to add redundancy and then distributing it across multiple disks. We determine the download time for two content access models – the fountain and fork-join models that involve simultaneous content access, and individual access from enqueued user requests respectively. For the fountain model we explicitly characterize the download time, while in the fork-join model we derive the upper and lower bounds. Our results show that coding reduces download time, through the diversity of distributing the data across more disks, even for the total storage used.

I. INTRODUCTION

Consumers of cloud storage and content centric networking demand that their content be reliably stored and quickly accessible. Cloud storage providers today strive to meet both demands by simply replicating content throughout the storage network over multiple disks. A large body of recent literature proposes erasure coding as a more efficient way to provide reliability.

Research in coding for distributed storage was galvanized by the results reported in [1]. Prior to that work, literature on distributed storage recognized that when compared with replication, coding can offer huge storage savings for the same reliability levels. But it was also argued that the benefits of coding are limited, and can easily be outweighed by certain disadvantages and extra complexity. Namely, to provide reliability in multi-disk storage systems, when some disks fail, it must be possible to restore either the exact lost data or an equivalent reliability with minimal download from the remaining storage. The cost of this repair regeneration was considered much higher in coded than in replication system [2], until [1] established existence and advantages of new regenerating codes. This work was then quickly followed, and the area is very active today (see e.g., [3], [4] and references therein).

A related line of work is concerned with another potential weakness of coding in distributed storage. Namely, if any part of the data changes, the corresponding coded packets must be updated accordingly. To minimize the cost of such updates, the authors in [5] propose a class of randomized codes which have update complexity scaling logarithmically with the size of data but can correct a linearly scaled number of disk failures. Furthermore, the existence of such update

efficient codes that also minimize the repair bandwidth for exact data reconstruction was established in [5].

Content accessibility is another main property of interest. In current multi-disk, cloud storage systems (e.g., Amazon), content files stored on the same disks may be simultaneously requested by multiple users. The file accessibility, therefore, depends on the dynamics of requests, and is limited by the disks' I/O bandwidth. In practice, it is again commonly improved by replicating content on multiple disks, which in turn requires more energy. Only recently was it realized that erasure coding can guarantee the same level of content accessibility with lower storage than replication. [6], [7]. In [7], it is considered that when there are multiple access requests, all but one of them are blocked, and the accessibility is measured in terms of blocking probability. In [6], multiple requests are placed in a queue instead of blocking. The authors propose a scheduling scheme to map requests to servers (or disks) to minimize the waiting time.

In this paper, we assume that requests that cannot be served upon arrival are queued, but we measure the accessibility in terms of the download time which includes the waiting time in queues and the time taken to read the data from the disk, which could be random. When the content is available redundantly on multiple disks, it is sufficient to read it only from a subset of these disks in order to recover the data. The key contribution of our work is to analyze how waiting for only a subset of disks to be read, provides diversity in storage which helps achieve a significant reduction in the download time.

Using redundancy in coding for delay reduction has also been studied in packet transmission [8]–[10], and in some other scenarios of content retrieval in [11]. Although they share some common spirit, they do not consider storage systems and the impact of redundancy coding in such scenarios.

This paper is organized as follows. In Section II we introduce the two content access models investigated in this work. In Section III we present the central idea of how coding gives diversity and reduces download time. Then we determine the download time for the two models in Section IV and Section V respectively. Finally we conclude in Section VI.

II. TWO CONTENT ACCESS/DELIVERY MODELS

We consider two specific content access models in order to isolate and emphasize different sources of delay in content

delivery, and show how they can be addressed through coding. In general, a content delivery model is some hybrid of the two models considered here.

A. The Fountain Model

Our first model can describe, e.g., a content delivery network scenario, where content may not be available at the point of request, and there is a delay associated with waiting for the content to become available at the contacted server. Once in possession of the content, the server can deliver it to the users by, e.g., broadcast enabled by digital fountain codes [12]. Thus multiple users do not affect each other's delivery time. Another scenario that can be modeled in this way is when content is broadcast at some prescribed times, but the arrival of users is random. We will refer to this model as the *fountain* model (where fountains are turned on at random times).

When a content request arrives at the server, the content has to be fetched from the distribution network and then transmitted to the user. The waiting time to obtain the content from the network is a non-negative random variable W . Once the content is obtained, the time to deliver it to the user is a positive random variable D , which is proportional to the size of content and erasure rate of the channel. Multiple users can access a server simultaneously and the content is broadcast to these users. Hence, the response time of the system (the download completion time) is $W + D$, and is independent of the number of requests being served simultaneously.

B. The Queueing Model

Our second model can describe, e.g., a storage area network (SAN), where content is stored on a disk, which can be accessed only by one user at a given time. The delay in this model is associated with the response time of the queueing systems. In this model, multiple requests by the same user do affect each other's content download time. We will refer to this model as the *queueing* model.

When a content request arrives at the disk, it enters a first-come-first-serve queue. After a request reaches the head of the queue, it takes some random service time to read the content from the disk. We model this service time a random variable with mean $1/\mu$. Here again, the download time is the sum of two components: the waiting time in queue and the service time required to read from the disk.

III. REDUCING DELAY BY CODING

In both of our models, the download completion time is a random variable. One natural way to reduce this time is to replicate the content across n independent servers (or disks). Then if the user issues n requests, one to each of the n servers, it only needs to wait for the one of the requests to be served. This strategy gives a sharp reduction in download time, but at the cost of n times more storage space and the cost of processing multiple requests.

We thus argue that it is more economical to divide the content into k blocks, encode them into $n > k$ coded blocks, and store them on n different servers (one block per server).

Each incoming request is sent to all n servers, and the content can be recovered when any k out of n blocks are successfully downloaded.

This can be achieved by using an (n, k) maximum distance separable (MDS) code to encode the k blocks of content. MDS codes have been suggested to provide reliability against server outages (or disk failures). In this paper we show that, in addition to error-correction, we can exploit these codes to reduce the download time of the content.

Note that for the fountain model, since multiple users can simultaneously access the content, the response times (waiting plus delivery time) of the n servers are independent. The download time is the k^{th} order statistic of the response time of each server. However, the analysis of download time for the queueing model is more challenging because the response times of the n queues are not independent.

Since in both models we require the first k out of n blocks to be downloaded, we now provide some background of finding the k^{th} order statistic of n independent and identically distributed (i.i.d) random variables. For a more complete treatment in order statistics, please refer to [13].

Let X_1, X_2, \dots, X_n be i.i.d. random variables. Then, $X_{k,n}$, the k^{th} order statistic of X_i , $1 \leq i \leq n$, or the k^{th} smallest variable has the distribution,

$$f_{X_{k,n}}(x) = n \binom{n-1}{k-1} F_X(x)^{k-1} (1 - F_X(x))^{n-k} f_X(x) \quad (1)$$

where $F_X(x)$ and $f_X(x)$ and the distribution and density functions of X respectively. In particular, if X_i 's are exponential with mean $1/\mu$, then the expectation and variance of order statistic $X_{k,n}$ are given by,

$$E[X_{k,n}] = \frac{1}{\mu} \sum_{i=1}^k \frac{1}{n-k+i} = \frac{1}{\mu} (H_n - H_{n-k}) \quad (2)$$

$$V[X_{k,n}] = \frac{1}{\mu^2} \sum_{i=1}^k \frac{1}{(n-k+i)^2} = \frac{1}{\mu^2} (H_{n^2} - H_{(n-k)^2}), \quad (3)$$

where H_n and H_{n^2} are generalized harmonic numbers defined as

$$H_n = \sum_{j=1}^n \frac{1}{j} \quad \text{and} \quad H_{n^2} = \sum_{j=1}^n \frac{1}{j^2}. \quad (4)$$

Note that $E[X_{k,n}]$ decreases as k decreases for a given n . This fact will provide us some intuition to understand the analysis of download time for the fountain and queueing models presented in Section IV and Section V respectively.

IV. MULTIPLE FOUNTAINS

In this section we investigate the redundancy storage in the context of fountain model. Content requests (e.g., request for videos, news or other information) from customers are sent to a network of servers. We focus in particular on *multiple fountain* content retrieval systems defined as follows.

Definition 1 ((n, k) multiple fountain): An (n, k) multiple fountain content retrieval system contains n servers.

Every content request entering the system is forked to n servers. Requests are served as soon as the content becomes available, which happens at a random time independently of request arrivals. A request is satisfied when any k out of n servers have responded and delivered their messages.

Recall that the content is divided into k blocks and encoded into $n > k$ coded blocks which are stored on n servers (one block per server). Content is said to be downloaded when any k out of n blocks are successfully delivered to the user. The fountain model described in Section II-A assumes that multiple users can access the server simultaneously (i.e., no queueing). The response time for each server is the sum of waiting time for the content to become available and the time taken to deliver each content block. We model the waiting time W as an exponentially distributed random variable with mean $1/\mu$. After the content becomes available, the server delivers it to the customer in constant time $\frac{D}{k}$, where the factor $1/k$ appears because each server only delivers $1/k$ fraction of the content.

The mean response time, i.e., the time taken to download the content, from the (n, k) multiple fountain system is given by the following theorem.

Theorem 1: The mean response time $T_{(n,k)}$ of a content retrieval system is

$$T_{(n,k)} = \frac{1}{\mu}(H_n - H_{n-k}) + \frac{D}{k}, \quad (5)$$

where H_n is defined in (4).

Proof: Since each message is delivered to the customer in constant time $\frac{D}{k}$, the mean response time for a request equals the waiting time until the content becomes available at k servers plus the delivery time $\frac{D}{k}$. The expected waiting time is the k^{th} order statistic of n i.i.d. exponential random variables with mean $1/\mu$, which has mean $\frac{1}{\mu}(H_n - H_{n-k})$ (c.f. (2)). ■

We notice that it is possible to have an optimal k such that (5) is minimized. The intuition behind this is the trade-off between the waiting time $\frac{1}{\mu}(H_n - H_{n-k})$ and the content delivery time $\frac{D}{k}$, as k varies from 1 to n . When k is small, the $T_{(n,k)}$ is dominated by the delivery time $\frac{D}{k}$. But as k increases, $T_{(n,k)}$ is dominated by the waiting time due to the increase in $\frac{1}{\mu}(H_n - H_{n-k})$, and decrease in $\frac{D}{k}$. The following lemma gives the optimal value of k .

Lemma 1: The k that minimizes (5) is given by

$$k^* = \arg \min_k T_{(n,k)} \approx \left\lceil \frac{-D\mu + \sqrt{D^2\mu^2 + 4nD\mu}}{2} \right\rceil. \quad (6)$$

Proof: We use the log approximation for H_n , i.e., $H_n \approx \log(n) + O(1)$ and $H_{n-k} \approx \log(n-k) + O(1)$. Substitute in (5) we obtain

$$T_{(n,k)} = \frac{1}{\mu} \log \left(\frac{n}{n-k} \right) + \frac{D}{k}.$$

Taking derivative with respect to k and set it to 0 we obtain

$$\frac{1}{\mu} k^2 + Dk - Dn = 0,$$

which has root $k^* = \frac{-D\mu + \sqrt{D^2\mu^2 + 4nD\mu}}{2}$ as in (6). ■

Fig. 1 shows the mean response time $T_{10,k}$ versus k for the $(10, k)$ multiple fountain system with parameters delivery time $D = 5$ and various values of mean waiting time $1/\mu$. We observe that the optimal value of k^* increases with $1/\mu$.




Fig. 1. Mean response time simulation for $(10, k)$ multiple fountain system with parameters $\frac{1}{\mu} = 0, 2, 4, 6$ and $D = 5$. Note that $1/\mu = 0$ means that the content is immediately available thus the download time decreases as k increases. On the contrary, when $D = 0$ (not shown in this plot), the download time goes down as k decreases.

V. FORK-JOIN QUEUES

In this section we consider the second content delivery model described in Section III, the queueing model. We show how coding can help minimize the time taken to download time of a content which is stored on an array of disks. We refer to this time as the response time. Although we focus on this storage model, it is possible to extend our results to other distributed systems such as parallel cluster computing [14].

A. System Model

Consider that a content F of unit size, divided into k blocks of equal size. It is encoded to $n > k$ blocks using a (n, k) maximum distance separable (MDS) code, and the coded blocks are stored on an array of n disks. MDS codes have the property that any k out of the n blocks are sufficient to reconstruct the entire file. An illustrative example with $n = 3$ disks and $k = 2$ is shown in Fig. 2. The content F is split into equal blocks a and b , and stored on 3 disks as a , b , and $a \oplus b$, the exclusive-or of blocks a and b . Thus each disk stores content of half the size of file F . Downloads from any 2 disks jointly enable reconstruction of F . Each user's request for content F is forked to all the n disks. Our objective is to determine the mean response time of the system – the expected time from the arrival of a request until it finishes service by reading the content from some k of the n disks. In order to evaluate the response time we model it as an n -fork k -join system which is defined as follows.

Definition 2 ((n, k) fork-join system): An (n, k) fork-join system consists of n processing nodes (fork nodes). Every




Fig. 2. A (3,2) fork-join system; storage is 50% higher, but response time (per disk & overall) is reduced.

arriving job is divided into n tasks which are sent to the queue at each of the n nodes. A task is served when it arrives at the top of its queue. The job departs the system when any k out of n tasks are served by their respective nodes. The remaining $n - k$ tasks abandon their queues and exit the system without receiving service.

The (n, n) fork-join system, known in literature as fork-join queue, has been extensively studied in, e.g., [15]–[17]. However, the (n, k) generalization in Definition 2 above has not been previously studied to our best knowledge.

We consider an (n, k) fork-join system where each node represents a disk from which content is being downloaded. Download requests arrive according to a Poisson process with rate λ per second. Every request is sent to each of the n disks. The time taken to download one unit of data is exponential with mean $1/\mu$. Since, each disk is requested to provide $1/k$ units of data, the service time for each node is exponentially distributed with mean $1/\mu'$ where $\mu' = k\mu$. Define the load factor $\rho \triangleq \lambda/\mu'$. We assume $\rho > 1$, or equivalently $\mu' > \lambda$ to ensure stability of the queue at each fork node.

B. Bounds on Mean Response Time

Our objective is to evaluate the mean response time $T_{(n,k)}$ of the (n, k) fork-join system described in Section V-A. It is the time from the arrival of a job until k out of n of its tasks are served by their respective nodes.

Even for the (n, n) system, the mean response time has not been found in closed form – only bounds are known. An exact expression for the response time is found only for the (2,2) fork-join in [16]. The reason why the fork-join system is harder to analyze than a set of parallel independent $M/M/1$ queues is that each incoming job is sent to the n queues. Hence, the arrivals to the queues are perfectly synchronized and the response times of the n queues are correlated.

The simplest case of an (n, k) fork-join system is the $(n, 1)$ system. It is not hard to see that this system behaves exactly as an $M/M/1$ queue with arrival rate λ and service rate $\mu' = n\mu$. Therefore its response time is exponential with the mean $T_{(n,1)}$ equal to $1/(n\mu - \lambda)$. It is difficult to evaluate $T_{(n,k)}$ exactly for other cases, but the bounds we derive below are fairly tight.

Theorem 2 (Upper Bound on Mean Response Time):

The mean response time $T_{(n,k)}$ of an (n, k) fork-join system satisfies

$$T_{(n,k)} \leq \frac{H_n - H_{n-k}}{\mu'} + \frac{\lambda[(H_{n^2} - H_{(n-k)^2}) + (H_n - H_{(n-k)})^2]}{2\mu'^2[1 - \rho(H_n - H_{n-k})]} \quad (7)$$

where λ is the request arrival rate, $\mu' = k\mu$ is the service rate at each queue, $\rho = \lambda/\mu'$ is the load factor, and the generalized harmonic numbers H_n and H_{n^2} are as given in (4).

Proof: We use a related, but easier to analyze queueing model called the split-merge system, to find this upper bound on $T_{(n,k)}$. In the (n, k) fork-join queueing model, after a node serves one of the tasks, it is free to process the next task in its queue. On the contrary, in the split-merge model, all n nodes are blocked until k of them finish service. Thus, the job departs all the queues at the same time. Since the nodes are not blocked in the fork-join system, the mean response time of the (n, k) split-merge model is an upper bound on (and a pessimistic estimate of) $T_{(n,k)}$ for the (n, k) fork-join system.

The (n, k) split-merge system is equivalent to an $M/G/1$ queue where arrivals are Poisson with rate λ and service time is a random variable S distribution according to the k^{th} order statistic of the exponential distribution. The mean and variance of S are (c.f. (2) and (3))

$$E[S] = \frac{H_n - H_{n-k}}{\mu'} \quad \text{and} \quad V[S] = \frac{H_{n^2} - H_{(n-k)^2}}{\mu'^2}. \quad (8)$$

The Pollaczek-Khinchin formula [18] gives the mean response time T of an $M/G/1$ queue in terms of the mean and variance of S as follows.

$$T = E[S] + \frac{\lambda E[S^2]}{2(1 - \lambda E[S])} \quad (9)$$

where the second moment $E[S^2] = V[S] + E[S]^2$. Substituting the values of $E[S]$ and $V[S]$ given by (8), we get the upper bound (7). ■

Remark 1: Note that the approach used in [16] to find an upper bound on the mean response time of the (n, n) fork-join system cannot be extended to the (n, k) fork-join system considered here. The authors in [16] show that the response times of the n queues form a set of associated random variables [19]. Associated random variables have the property that their expected maximum is less than that for independent variables with the same marginal distributions. Thus, the mean response time of the (n, n) fork-join system is upper bounded by that of the system of n independent $M/M/1$ queues. However, this property of associated variables does not hold for the k^{th} order statistic for $k < n$.

Theorem 3 (Lower Bound on Mean Response Time):

The mean response time $T_{(n,k)}$ of an (n, k) fork-join queueing system satisfies

$$T_{(n,k)} \geq \frac{1}{\mu'} [H_n - H_{n-k} + \rho(H_{n(n-\rho)} - H_{(n-k)(n-k-\rho)})] \quad (10)$$

where λ is the request arrival rate, $\mu' = k\mu$ is the service rate at each queue, $\rho = \lambda/\mu'$ is the load factor, and the generalized harmonic number $H_{n(n-\rho)}$ is given by

$$H_{n(n-\rho)} = \sum_{j=1}^n \frac{1}{j(j-\rho)}.$$

Proof: The lower bound in (10) is a generalization of the bound for the (n, n) fork-join system derived in [17]. The bound for the (n, n) system is derived by considering that a job goes through n stages of processing. A job is said to be in the j^{th} stage, for $0 \leq j \leq n-1$, if j out of n tasks have been served by their respective nodes and the remaining $n-j$ tasks are waiting to be served. The job will depart the system when all n tasks are served.

For the (n, k) fork-join system, since we only need k tasks to finish service, the number of stages of processing is reduced. Each job now goes through k stages of processing, where in the j^{th} stage, for $0 \leq j \leq k-1$, j tasks have finished processing and we are waiting for the $k-j$ more tasks to finish service in order to complete the job.

Consider two jobs B_1 and B_2 in the i^{th} and j^{th} stages of processing respectively. Let $i > j$, or in other words, B_1 has completed more tasks than B_2 . Since every incoming job is sent to all n queues, this implies that B_1 's tasks will be in front B_2 's in all $n-i$ queues remaining to be served for B_1 . Further, we can conclude that the mean service rate of job B_2 moving to the $(j+1)^{\text{th}}$ stage of processing is at most $(n-j)\mu'$. If the $n-j$ pending tasks are at the head of all the respective queues, then the service rate will be exactly $(n-j)\mu'$. However, B_1 's task could be ahead of B_2 's in one of the $n-j$ pending queues, due to which that task of B_2 cannot be immediately served. Hence, we have shown that for a job in the j^{th} stage of processing, the mean service rate is at most $(n-j)\mu'$.

Consider an $M/M/1$ queue with arrival rate λ and service rate $(n-j)\mu'$. Its response time is exponentially distributed with mean $T_j = 1/((n-j)\mu' - \lambda)$. By the memoryless property of the exponential distribution, the total mean response time is the sum of the mean response times of each of the k stages of processing, given by

$$\begin{aligned} T_{(n,k)} &\geq \sum_{j=0}^{k-1} \frac{1}{(n-j)\mu' - \lambda} = \frac{1}{\mu'} \sum_{j=0}^{k-1} \frac{1}{(n-j) - \rho} \\ &= \frac{1}{\mu'} \sum_{j=0}^{k-1} \left[\frac{1}{n-j} + \frac{\rho}{(n-j)(n-j-\rho)} \right] \\ &= \frac{1}{\mu'} [H_n - H_{n-k} + \rho \cdot (H_{n(n-\rho)} - H_{(n-k)(n-k-\rho)})] \end{aligned}$$

■

Hence, we have found lower and upper bounds on the mean response time $T_{(n,k)}$. In Section V-D, we perform simulations to check the tightness of these bounds. These help us answer some practical questions in designing storage systems with the minimum download completion time.

C. Extension to General Service Time Distribution

In this section we derive the upper bound on expected download time with a general service time distribution at each node, instead of the exponential service time considered so far. Let X_1, X_2, \dots, X_n be the i.i.d random variables representing the service times of the n nodes, with expectation $E[X_i] = \frac{1}{\mu'}$ and variance $V[X_i] = \sigma^2$ for all i .

Theorem 4 (Upper Bound with General Service Time):

The mean response time $T_{n,k}$ of an (n, k) fork-join system with general service time X such that $E[X] = \frac{1}{\mu'}$ and $V[X] = \sigma^2$ satisfies

$$\begin{aligned} T_{(n,k)} &\leq \frac{1}{\mu'} + \sigma \sqrt{\frac{k-1}{n-k+1}} \\ &\quad + \frac{\lambda \left[\left(\frac{1}{\mu'} + \sigma \sqrt{\frac{k-1}{n-k+1}} \right)^2 + \sigma^2 C(n, k) \right]}{2 \left[1 - \lambda \left(\frac{1}{\mu'} + \sigma \sqrt{\frac{k-1}{n-k+1}} \right) \right]}. \end{aligned} \quad (11)$$

Proof: The proof follows from Theorem 2 where the upper bound can be calculated using (n, k) split-merge system and Pollaczek-Khinchin formula (9). Unlike the exponential distribution, we do not have an exact expression for S , i.e., the k^{th} order statistic of the service times X_1, X_2, \dots, X_n . Instead, we use the following upper bounds on the expectation and variance of S derived in [20] and [21].

$$E[S] \leq \frac{1}{\mu'} + \sigma \sqrt{\frac{k-1}{n-k+1}} \quad (12)$$

$$V[S] \leq C(n, k)\sigma^2, \quad (13)$$

The proof of (12) involves Jensen's inequality and Cauchy-Schwarz inequality. For details please refer to [20]. The constant $C(n, k)$ depends only on n and k , and can be found in the table in [21]. Holding n constant, $C(n, k)$ decreases as k increases. The proof of (13) can be found in [21].

Note that (9) strictly increases as either $E[S]$ or $V[S]$ increases. Thus, we can substitute the upper bounds in it to obtain the upper bound on mean response time (11).

Finally, we note that our proof in Theorem 3 cannot be extended to this general service time setting. The proof requires memoryless property of the service time, which does not necessary hold in the general service time case. ■

D. Numerical Examples and Simulation

In this section we present numerical and simulation example results to help us appreciate how storing the content on n disks using an (n, k) fork-join system (as described in Section V-A) reduces the expected download time. The results demonstrate the tightness of the bounds derived in Section V-B. In addition, we simulate the fork-join system to obtain an empirical cumulative density function (CDF) for the download time.

The download time of a file with k blocks can be improved by increasing 1) the storage expansion n/k per file and/or 2) the number n of disks used for file storage. For example, fork-join systems $(4, 2)$ and $(10, 5)$ both provide a storage expansion of 2, but the former uses 4 and the latter 10 disks,

and thus their download times behave differently. Both the total storage and the number of storing elements could be a limiting factor in practice.

We first address the scenario where the number of disks n is kept constant, but the storage expansion changes from 1 to n as we choose k from n to 1. We then study the scenario where the storage expansion factor n/k is kept constant, but the number of disks varies.

1) *Flexible Storage Expansion & Fixed Number of Disks:*

In Fig. 3 we plot the mean response time $T_{(n,k)}$ versus k for fixed number of disks $n = 10$, arrival rate $\lambda = 1$ request per second and service rate $\mu = 3$ units of data per second. Each disk stores $1/k$ units of data and thus the service rate of each individual queue is $\mu' = k\mu$. The simulation plot




Fig. 3. Mean response time $T_{(n,k)}$ increases with k for fixed n because the redundancy of coding reduces. The plot also demonstrates the tightness of the bounds derived in Section V-B

shows that as k increases with n fixed, the code rate k/n increases thus reducing the amount of redundancy. Hence, $T_{(n,k)}$ increases with k . We also observe that the bounds (7) and (10) derived in Section V-B are very tight.

In addition to low mean response time, ensuring quality-of-service to the user may also require that the probability of exceeding some maximum tolerable response time is small. Thus, we study the CDF of the response time for different values of k for a fixed n .

In Fig. 4 we plot the CDF of the response time with $k = 1, 2, 5, 10$ for fixed $n = 10$. The arrival rate and service rate are $\lambda = 1$ and $\mu = 3$ as defined earlier. For $k = 1$, the PDF is represents the minimum of n exponential random variables, which is also exponentially distributed.

The CDF plot can be used to design a storage system that gives probabilistic bounds on the response time. For example, if we wish to keep the response time below 0.1 seconds with probability at least 0.75, then the CDF plot shows that $k = 5, 10$ satisfy this requirement but $k = 1$ does not. The plot also shows that at 0.4 seconds, 100% of requests are complete in all fork-join systems, but only 50% are complete in the single-disk case





Fig. 4. CDFs of the response time of $(10, k)$ fork-join systems, and the required storage

2) *Flexible Number of Disks & Fixed Storage Expansion*

: Now we take a different viewpoint and analyze the benefit of spreading the content across more disks while using the same total storage space. Fig. 5 shows a simulation plot of the mean response time $T_{(n,k)}$ versus k while keeping constant code rate $k/n = 1/2$. The response time $T_{(n,k)}$ reduces with increase in k because we get the diversity advantage of having more disks. With a very large n , as k increases, the theoretical bounds (7) and (10) suggest that $T_{(n,k)}$ approaches zero. This is because we assumed that service rate of a single disk $\mu' = k\mu$ since the $1/k$ units of the content F is stored on one disk. However, in practice the mean service time $1/\mu'$ will not go zero because reading each disk will need some non-zero processing delay in completing each task irrespective of the amount of data stored on it.

In order to understand the response time better, we plot its CDF in Fig. 6 for different values of k for fixed ratio $k/n = 1/2$. Again we observe that the diversity of increasing number of disks n helps reduce the response time.

VI. CONCLUSION AND FUTURE WORK

We analyzed the download time of a content file from a distributed storage system. We assume that content of interest is available redundantly on multiple disks, or on multiple




Fig. 5. The mean response time with constant code rate

nodes throughout the network, entirely or in chunks. Such scenarios may be a consequence of caching throughout a network or as a result of purposeful storage in data centers and storage area networks. Our idea is to also make redundant requests for content in order to reduce the download time through route diversity (the fountain model) and load balancing (the queuing model). Under this central idea, we showed that the expected download time is significantly reduced using coding – we divide the content into k parts, apply an (n, k) MDS code and store it on n disks. The file can be recovered by reading any k of the n disks. We analytically studied the mean response time and derived tight upper and lower bounds.

In practical storage systems, adding redundancy in storage not only requires extra capital investment in storage devices, networking and management but also consumes more energy. It has been estimated that around 40% of total operation cost has been related to power distribution, cooling, and electricity bills [22] and the total data center power consumption in 2005 was already 1% of the total US power consumption [23]. It would be interesting to study the fundamental tradeoff between power consumption and quality of service (QoS) performance and distill insight on system design. As we have shown in this paper, for the same performance, coding requires less redundancy than conventional replication based storage. We would like to investigate how much energy can coding based storage save us. Furthermore, this also motivates the research on more efficient load balancing algorithms, which not only fork each job onto a set of servers, but do so with power conservation in mind.

In this paper we do not consider some other possible costs, such as the decoding time required to reconstruct the original content out of k received blocks, or placing redundant requests. We try to qualitatively illustrate the possible benefits of coding without exactly quantifying the gains in particular, practical systems. Taking the decoding time (which affects delay performance) into consideration





Fig. 6. CDFs of the response time of $(2k, k)$ fork-join systems, and the required storage

motivates us to investigate the optimal redundancy level. We also leave this as our future work.

REFERENCES

- [1] A. G. Dimakis, P. B. Godfrey, M. Wainwright, and K. Ramchandran, "Network coding for distributed storage systems," in *Proc. of IEEE INFOCOM'07, Anchorage, AK, USA*, May, 2007, pp. 2000–2008.
- [2] R. Rodrigues and B. Liskov, "High availability in DHTs: Erasure coding vs. replication," in *Int. Workshop P. to P. Sys.*, 2005, pp. 226–239.
- [3] K. V. Rashmi, N. B. Shah, P. V. Kumar, and K. Ramchandran, "Explicit construction of optimal exact regenerating codes for distributed storage," *Allerton Conf. on Commun. Control and Comput.*, pp. 1243–1249, Sep. 2009.
- [4] N. B. Shah, K. V. Rashmi, P. V. Kumar, and K. Ramchandran, "Interference alignment in regenerating codes for distributed storage: necessity and code constructions," *IEEE Trans. Inform. Theory*, vol. 58, pp. 2134–2158, Apr. 2012.
- [5] A. S. Rawat, S. Vishwanath, A. Bhowmick, and E. Soljanin, "Update efficient codes for distributed storage," in *Proc. Int. Symp. Inform. Theory*, 2011, pp. 1457–1461.
- [6] L. Huang, S. Pawar, H. Zhang and Kannan Ramchandran, "Codes can reduce queuing delay in data centers," *Proc. Int. Symp. Inform. Theory*, pp. 2766–2770, Jul. 2012.
- [7] U. Fehrer, M. Médard, and E. Soljanin, "Toward sustainable networking: Storage area networks with network coding," *Allerton Conf. on Commun. Control and Comput.*, Oct. 2012.
- [8] G. Kabatiansky, E. Krouk, and S. Semenov, *Error correcting coding and security for data networks: analysis of the superchannel concept*, 1st ed. Wiley, Mar. 2005, ch. 7.

- [9] N. F. Maxemchuk, "Dispersity routing," *Proc. Int. Conf. Commun.*, pp. 41.10 – 41.13, Jun. 1975.
- [10] Y. Liu, J. Yang, and S. C. Draper, "Exploiting route diversity in multi-packet transmission using mutual information accumulation," *Allerton Conf. on Commun. Control and Comput.*, pp. 1793–1800, Sep. 2011.
- [11] E. Soljanin, "Reducing delay with coding in (mobile) multi-agent information transfer," *Allerton Conf. on Commun. Control and Comput.*, pp. 1428–1433, Sep. 2010.
- [12] A. Shokrollahi, "Raptor codes," *IEEE/ACM Trans. on Network.*, vol. 14, pp. 2551–2567, Jun. 2006.
- [13] S. Ross, *A first course in probability*, 6th ed. Prentice Hall, 2002, ch. 6.6, p. 273.
- [14] J. Dean and S. Ghemawat, "MapReduce: simplified data processing on large clusters," *ACM Commun. Mag.*, vol. 51, no. 1, pp. 107–113, Jan. 2008.
- [15] C. Kim and A. K. Agrawala, "Analysis of the fork-join queue," *IEEE Trans. Comput.*, vol. 38, no. 2, pp. 250–255, Feb. 1989.
- [16] R. Nelson and A. Tantawi, "Approximate analysis of fork/join synchronization in parallel queues," *IEEE Trans. Comput.*, vol. 37, no. 6, pp. 739–743, Jun. 1988.
- [17] A. M. E. Varki and H. Chen, "The M/M/1 fork-join queue with variable sub-tasks."
- [18] H. C. Tijms, *A first course in stochastic models*, 2nd ed. Wiley, 2003, ch. 2.5, p. 58.
- [19] J. Esary, F. Proschan and D. Walkup, "Association of random variables, with applications," *Annals of Math. Stat.*, vol. 38, no. 5, pp. 1466–1474, Oct. 1967.
- [20] B. C. Arnold and R. A. Groeneveld, "Bounds on expectations of linear systematic statistics based on dependent samples," *Annals of Stat.*, vol. 7, pp. 220–223, Oct. 1979.
- [21] N. Papadatos, "Maximum variance of order statistics," *Ann. Inst. Statist. Math.*, vol. 47, pp. 185–193, 1995.
- [22] A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel, "The cost of a cloud: research problems in data center networks," *SIGCOMM Comput. Commun. Rev.*, vol. 39, pp. 68–73, Dec. 2008.
- [23] V. Mathew, R. K. Sitaraman, and P. Shenoy, "Energy-aware load balancing in content delivery networks," *Proc. IEEE INFOCOM*, to appear, 2012.