

Memory Allocation in Distributed Storage Networks

Mohsen Sardari[†], Ricardo Restrepo[‡], Faramarz Fekri[†], Emina Soljanin^{*}

[†]School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332

[‡]School of Mathematics, Georgia Institute of Technology, Atlanta, GA 30332

^{*}Alcatel-Lucent Bell-Labs, Murray Hill, NJ 07974

Email: [†]{mohsen.sardari, fekri}@ece.gatech.edu, [‡]restrepo@math.gatech.edu, ^{*}emina@research.bell-labs.com

Abstract—We consider the problem of distributing a file in a network of storage nodes whose storage budget is limited but at least equals the size file. We first generate T encoded symbols (from the file) which are then distributed among the nodes. We investigate the optimal allocation of T encoded packets to the storage nodes such that the probability of reconstructing the file by using any r out of n nodes is maximized. Since the optimal allocation of encoded packets is difficult to find in general, we find another objective function which well approximates the original problem and yet is easier to optimize. We find the optimal symmetric allocation for all coding redundancy constraints using the equivalent approximate problem. We also investigate the optimal allocation in random graphs. Finally, we provide simulations to verify the theoretical results.

I. INTRODUCTION

A file in a distributed storage network can be replicated throughout the network to improve the performance of retrieval process, measured by routing efficiency, persistence of the file in the network when some storage locations go out of service, and many other criteria. Most of the studies in network file storage consider a common practice where every node in the network either stores the entire file or none of it. In an important article, Naor and Roth [1] studied how to store a file in a network such that every node can recover the file by accessing only the portions of the file stored on itself and its neighbors, with the objective of minimizing the total amount of data stored. By applying MDS (Maximum Distance Separable) codes and generating codeword symbols of the file, they presented a solution that is asymptotically optimal in minimizing the total number of stored bits, when the original file has a length much larger than the logarithm of the graph's degree of the storage network. Other works [2], [3] extended the result of [1] and devised algorithms for memory allocation in tree networks with heterogeneous clients. Distributed storage is also studied in sensor networks [4], [5]. In sensor networks, the focus is usually on the data retrieval assuming that a data collector has access to a random subset of storage nodes while in this paper we address the allocation problem.

One of the appealing features for a distributed storage system is the ability to scale the persistence of data arbitrarily up and down on-demand. In other words, the cost of accessing the stored data should be adjustable based on the demand. In one extreme, all the nodes have “easy” access to the stored file, either by storing the whole file or a large part of it. On the other extreme, just a single node stores the file entirely and other nodes need to fetch the file from that node. It is clear that

by making more copies of a file and spreading those copies in the network, the retrieval of the file becomes easier. The use of MDS codes provides the flexibility to increase the persistence of a file gradually. For example, for a given file of size F , we can generate T symbols using a (T, F) MDS code such that every F -subset of those T symbols is sufficient to reconstruct the original file. We call T the budget considered for the file. Now, the question is as to how increasing the budget of a file affects the retrieval process. In order to answer this question, we need to consider a model for data retrieval. Recently, Leong et. al. [6] investigated this problem and introduced the following model for the network. Consider a network with n storage nodes. We distribute a file of size F and budget T (packets or symbols) among these storage nodes. Then, we look at all the possible subsets of size r of the storage nodes. We say that a specific r -subset is successful in recovering the file if the total number of packets stored in that subset of the nodes is at least the file size F . We are to find the best assignment of these T symbols to n storage nodes such that the maximum number of the r -subsets of storage nodes have enough number of symbols to reconstruct the file. The rationale behind the model is that in a real storage network, every node can be reached by all the other nodes in network. Once a retrieval request for a file is received by a node in network, the node tries to fetch all the parts of the file and respond to the request. The cost of fetching the parts from different nodes is not equal (other nodes may be down, busy, etc.). Therefore, in the model we assume that each node fetches the necessary parts of the file from the other $r - 1$ most accessible nodes.

In general, this problem is quite challenging and the optimal allocation is non-trivial. In [6], the authors provide some results for the symmetric allocation and probability-1 recovery regime which is a special case of the problem introduced in [1]. Symmetric allocation refers to a scheme where, based on the budget, we split the storage nodes into two groups: the nodes with no stored symbols and the nodes that store the same number of symbols. In probability-1 recovery regime, all the nodes should be able to reconstruct the file. As illustrated in [6], the optimal allocation is not obvious even if we only consider the symmetric allocations.

For very low budgets, we observe that the budget is concentrated over a minimal subset of storage nodes in the optimal allocation. On the other hand, for high budget levels, we observe a maximal spread of budget over storage nodes. It is of interest to determine as to how this transition occurs and

also to study the behavior of the optimal allocation versus budget. In this paper, we take the initial steps towards the characterization of the optimal allocation. In section II, we give the formal definition of the problem and the model we consider. Then, in Section III we prove that an easier to solve problem well approximates the original problem. Using the alternative approach, we solve the file allocation problem for symmetric allocations (Section IV); we also consider symmetric allocations in random graphs. Finally, simulation results are provided in Section V.

II. FILE ALLOCATION PROBLEM

A. Problem Statement

We are given a file of size F and a network with budget T . We generate T redundant symbols using a (T, F) MDS code. An allocation of T symbols to n nodes is defined to be a partition of T into n sets of sizes x_1, \dots, x_n , where x_i is the number of symbols allocated to the i th storage node. Note that $\sum x_i = T$ and $x_i \geq 0$ for $i = 1, \dots, n$. Our goal is to find an allocation which maximizes the number of r -subsets jointly storing F or more packets.

Combination networks provide a simple illustration of the allocation problem under study. As shown in Figure 1, there




Fig. 1. Combination Network. Virtual source node in layer one has a file of size F . The solid nodes in layer two represent the storage nodes in network. The third layer contains virtual receiver nodes. Each receiver node corresponds to an r -subset of the storage nodes.

are three layers of nodes. A virtual source node in layer 1 has a file of size F (packets) to be distributed among storage nodes in layer 2. There are n storage nodes in layer two which represent the actual storage nodes in the storage network. Attributed to the file is a budget T . The data retrieval phase is visualized in the third layer of the combination network, which contains $\binom{n}{r}$ virtual receiver nodes. Each receiver node corresponds to an r -subset of the storage nodes. We are going to find the best allocation of the budget T such that the maximum number of receivers R in the combination network can reconstruct the file. Indeed, the success of the recovery process depends on the budget T . Based on the illustration in Figure 1, we use the terms receiver and r -subset interchangeably.

We use these notations throughout the paper:

- $[m] = \{1, \dots, m\}$ and $[m]_* = \{0, 1, \dots, m\}$

- $A^r = \{(s_1, \dots, s_r) : s_i \in A\}$. Note that there is no limit on the number of times an element s_k in set A can be chosen in (s_1, \dots, s_r) .
- $A^{[r]} = \{(s_1, \dots, s_r) : s_i \in A \text{ and } s_i \neq s_j \text{ for } i \neq j\}$. In other words, $A^{[r]}$ is the set of ordered vectors with distinct elements.
- $d_{<F}^u(\cdot)$ is an operator on polynomials which truncates to the terms of degree less than F with respect to u .

Furthermore, we use the notation \mathbb{I} for indicator function, defined as

$$\mathbb{I}_{(\omega \in \Omega)} = \begin{cases} 1 & \text{if } \omega \in \Omega \\ 0 & \text{if } \omega \notin \Omega \end{cases}.$$

For an allocation (x_1, \dots, x_n) of T symbols, let $\Psi(x_1, \dots, x_n)$ count the number of unsuccessful receivers. We can write Ψ as

$$\Psi(x_1, \dots, x_n) := \sum_{\substack{S \subseteq [n] \\ |S|=r}} \mathbb{I}\left(\sum_{i \in S} x_i < F\right),$$

where the first sum is over all the subsets of size r of storage nodes. Therefore, the allocation problem we consider is the following optimization problem:

$$\begin{aligned} & \text{minimize} && \Psi(x_1, \dots, x_n) \\ & \text{subject to} && \sum_{i=1}^n x_i = T \\ & && x_i \geq 0, \quad x_i \text{ integer} \end{aligned}.$$

It is challenging to find the optimal allocation because of the large space of possible allocations, non-convexity, and discontinuity of the indicator function. Our approach for solving this problem is to look into another quantity which, for $r \ll \sqrt{n}$, closely approximates Ψ but it is easier to compute.

B. Main Result

Let α_k be the fraction of nodes containing k symbols, and let $c := T/n$. The set of constraints on admissible allocation with respect to α can be re-written as

$$\left\{ \alpha : \sum_{i=0}^F \alpha_i = 1 \text{ and } \sum_{i=0}^F i \alpha_i = c \right\}.$$

Given an allocation (x_1, \dots, x_n) , we can compute the parameters $\alpha_0, \dots, \alpha_F$. Then, we define $\varphi(\alpha_0, \dots, \alpha_F)$ as the probability that a receiver with access to a uniformly chosen subset of nodes s from $[n]^r$ (shown by $s \sim [n]^r$) is unsuccessful in recovering the file. We have

$$\varphi(\alpha_0, \dots, \alpha_F) := \mathbf{P}_{s \sim [n]^r} \left(\sum_{i=1}^r x_{s_i} < F \mid \alpha \right). \quad (1)$$

Our first claim says that φ is a good approximation for Ψ .

Theorem 1:

$$\inf_{\alpha} \varphi(\alpha) - \frac{2(r-1)^2}{n} \leq \frac{r!}{n^r} \inf_{x \in \text{Allocations}} \Psi(x) \leq \inf_{\alpha} \varphi(\alpha).$$

Proof: Proof is given in Sec. III. ■

Dealing with the functional φ is simpler than working with Ψ . In the definition of φ , the random vector is chosen from $[n]^r$ where repetition is allowed. As a result, the probability

generating function of φ has a simple form and is easy to work with. The main result of the Theorem 1 is that for $r \ll \sqrt{n}$, we can solve the problem of minimizing $\varphi(\alpha_0, \dots, \alpha_F)$ instead, which is simpler than solving for the original optimization problem. Moreover, this solution is also a good approximation of the problem. We will further discuss the discrepancy in the optimal solution through an example in the last section.

From this point on, we will drop the conditioning on α for brevity. Please note that φ is just a function of $(\alpha_0, \dots, \alpha_F)$ and its value remains the same for all allocations with the same $(\alpha_0, \dots, \alpha_F)$.

III. DISCUSSIONS AND PROOF OF THE MAIN RESULT

Consider a receiver which has access to the vector of storage nodes $\mathbf{s} = [s_1, \dots, s_r]$, where \mathbf{s} is uniformly chosen from $[n]^{[r]}$. Let $\mathbf{P}_{\mathbf{s} \sim [n]^{[r]}} [\sum_{i=1}^r x_{s_i} < F]$ represent the probability that the total number of symbols stored in a randomly chosen set of size r of storage nodes \mathbf{s} is less than the file size F . There are in total $n(n-1)\dots(n-(r-1)) = r! \binom{n}{r}$ ordered vectors like \mathbf{s} in $[n]^{[r]}$ (note that the subsets in $[n]^{[r]}$ are ordered). Although working with ordered sets is slightly more complicated, as we will see shortly, this will help us in finding a better approximation for Ψ .

The total number of unsuccessful receivers $\Psi(x_1, \dots, x_n)$ can be calculated easily if we have the probability $\mathbf{P}_{\mathbf{s} \sim [n]^{[r]}}$ that a receiver with access to a randomly chosen subset of nodes \mathbf{s} is unsuccessful in recovering the file. In the definition of Ψ , we are only concerned with the total number of unsuccessful receivers. If we choose \mathbf{s} from a space like $[n]^{[r]}$ where order is important, we need to eliminate the effect of over-counting. Here, since $\mathbf{s} \sim [n]^{[r]}$, a division by $r!$ is sufficient. Hence, the functional $\Psi(x_1, \dots, x_n)$ can be re-written as

$$\Psi(x_1, \dots, x_n) = \binom{n}{r} \mathbf{P}_{\mathbf{s} \sim [n]^{[r]}} \left[\sum_{i=1}^r x_{s_i} < F \right]. \quad (2)$$

In order to prove Theorem 1, we first derive the lower bound on φ and then we prove the upper bound in the lemmas below.

Lemma 1: For any allocation (x_1, \dots, x_n) , satisfying a given set of α 's, the following hold:

$$\frac{r! \Psi(x_1, \dots, x_n)}{n^r} \leq \varphi(\alpha_0, \dots, \alpha_F). \quad (3)$$

Proof of Lemma 1: The inequality (3) follows immediately from the definitions of φ and Ψ ((1) and (2)). ■

In order to prove the upper bound, we need to look at the total variation between the distributions of a uniform random vector $\mathbf{s} \sim [n]^{[r]}$ and a uniform random vector $\mathbf{s}' \sim [n]^{[r]}$.

Definition 1: The total variation of two probability distributions μ and ν on a discrete space Ω is defined as

$$\text{TV}(\mu, \nu) = \sup_{A \subseteq \Omega} |\mu(A) - \nu(A)|.$$

A well known integral formula for the total variation between two distributions is given by

$$\text{TV}(\mu, \nu) = \frac{1}{2} \sum_{\omega \in \Omega} |\mu(\omega) - \nu(\omega)|.$$

Lemma 2: Let $\Omega = [n]^r$. Further, let μ be the uniform probability distribution over Ω , and ν be the uniform probability distribution over the subset S of Ω consisting of vectors with distinct entries; ν is 0 on $\Omega \setminus S$. Then, we have

$$\text{TV}(\mu, \nu) \leq \frac{(r-1)^2}{n}.$$

Proof of Lemma 2: The total number of non-repetitive vectors of size r in Ω is $n(n-1)\dots(n-r+1)$. We use the short hand $n^{[r]}$ for this expression. Then we can write

$$\begin{aligned} \text{TV}(\mu, \nu) &= \frac{1}{2} \left\{ \sum_{\omega \in S} |\mu(\omega) - \nu(\omega)| + \sum_{\omega \notin S} |\mu(\omega) - \nu(\omega)| \right\} \\ &= \frac{1}{2} \left\{ \sum_{\omega \in S} \left(\frac{1}{n^{[r]}} - \frac{1}{n^r} \right) + \sum_{\omega \notin S} \frac{1}{n^r} \right\} \\ &= 1 - \frac{n^{[r]}}{n^r} = 1 - \prod_{l=0}^{r-1} \left(1 - \frac{l}{n} \right) < 1 - \left(1 - \frac{r-1}{n} \right)^{r-1} \\ &< 1 - \left[1 - (r-1) \cdot \frac{r-1}{n} \right] = \frac{(r-1)^2}{n}. \end{aligned}$$

Lemma 3:

$$\frac{r! \Psi(x_1, \dots, x_n)}{n^r} \geq \varphi(\alpha_0, \dots, \alpha_F) - \frac{2(r-1)^2}{n} \quad (4)$$

Proof of Lemma 3: Using the results of Lemma 2 and definitions of Ψ and φ , we can write

$$\begin{aligned} & \left| \frac{r! \Psi(x_1, \dots, x_n)}{n^r} - \varphi(\alpha_0, \dots, \alpha_F) \right| \\ &= \left| \frac{n^{[r]}}{n^r} \mathbf{P}_{\mathbf{s} \sim [n]^{[r]}} \left[\sum_{i=1}^r x_{s_i} < F \right] - \mathbf{P}_{\mathbf{s} \sim [n]^r} \left(\sum_{i=1}^r x_{s_i} < F \right) \right| \\ &\leq \left| \mathbf{P}_{\mathbf{s} \sim [n]^{[r]}} \left[\sum_{i=1}^r x_{s_i} < F \right] - \mathbf{P}_{\mathbf{s} \sim [n]^r} \left(\sum_{i=1}^r x_{s_i} < F \right) \right| \\ &\quad + \left(1 - \frac{n^{[r]}}{n^r} \right) \leq \frac{2(r-1)^2}{n}. \end{aligned}$$

The first inequality above follows from the triangle inequality, and the second from Lemma 2. ■

The proof of the Theorem 1 follows from Lemma 1 and 3.

IV. OPTIMAL SYMMETRIC ALLOCATIONS

Following the results of the previous section, for cases where $r \ll \sqrt{n}$, we have $\frac{2(r-1)^2}{n} \ll 1$ and therefore, finding the optimal allocation of the symbols is equivalent to minimizing the function $\varphi(\alpha_0, \dots, \alpha_F)$. In this section, we direct our attention to symmetric allocations. In the case of symmetric allocations, we can find the optimal symmetric allocation and probability of success for all different budgets T . An allocation is called symmetric if we allocate the budget T as follows: we pick a number, say j , and we allocate chunks of size T/j until we run out of the budget. Now, we have two types of nodes: fraction α_0 of nodes which are left empty and the fraction α_j of the nodes which store j number of symbols.

Again, the optimal allocation is not obvious even if we consider only symmetric allocations. For instance, for very low budgets ($T \approx F$), we can easily argue that the budget should be concentrated over a minimal subset of nodes. For example, consider the case where $T = F$, if we store the entire file over one of the storage nodes, then the total number of successful receivers is $\binom{n-1}{r-1}$. If we break the file into two parts each of size $F/2$, then the total number of successful receivers is going to be $\binom{n-2}{r-2}$. By using the well-known identity

$$\binom{n}{r} = \binom{n-1}{r-1} + \binom{n-1}{r},$$

it is clear that the former allocation outperforms the latter. Similarly, other symmetric allocations can also be rejected. When the budget is very high ($T \approx nF/r$), the budget should be spread maximally. For example, consider the case where $T = nF/r$. In this case, by spreading the budget over all the storage nodes, we can achieve the probability-1 recovery. If one distributes this budget by allocating chunks of size F (storing the file in its entirety), he will be worse-off since the probability of success will be

$$1 - \binom{n - \lfloor n/r \rfloor}{r},$$

which is clearly less than 1. This behavior gives rise to questions like: ‘‘When to switch from minimal to maximal spread of the budget?’’, ‘‘Is there any situation where there exists a solution other than minimal or maximal spreading?’’

First, we give a useful expression for φ in the lemma below, which is simpler to work with. Then, we investigate the optimal symmetric allocation.

Lemma 4:

$$\varphi(\alpha_0, \dots, \alpha_F) = \left[\mathbf{d}_{<F}^u \left(\sum_{k=0}^F u^k \alpha_k \right)^r \right]_{|u=1} \quad (5)$$

Proof of Lemma 4: If s_i is a random element of $[n]$, then the probability that $\mathbf{P}(x_{s_i} = k)$ is equal to α_k . Therefore, the probability generating function of x_{s_i} is equal to $\sum_{k=0}^F u^k \alpha_k$. Hence, if $s = (s_1, \dots, s_r)$ is a uniform random vector in $[n]^r$, then the probability generating function of $\sum_{i=1}^r x_{s_i}$ is equal to $(\sum_{k=0}^F u^k \alpha_k)^r$. It follows then that

$$\mathbf{P}_{s \sim [n]^r} \left(\sum_{i=1}^r x_{s_i} < F \right) = \left[\mathbf{d}_{<F}^u \left(\sum_{k=0}^F u^k \alpha_k \right)^r \right]_{|u=1}. \quad (6)$$

and (5) is immediate. \blacksquare

In a symmetric allocation, suppose that the fraction of the non-empty nodes is α_j with j number of symbols each. Therefore, in the expression of $\varphi(\alpha_0, \dots, \alpha_F)$ at most α_0 and α_j have non-zero values. Our goal is to find the optimal value of j .

In this case, using Lemma 4, the problem of minimizing $\varphi(\alpha_0, \dots, \alpha_F)$ over $\{\alpha_0 + \alpha_j = 1, j\alpha_j = c\}$ reduces to

$$\varphi(\alpha_j) = \left[\mathbf{d}_{<F}^u (\alpha_0 + u^j \alpha_j)^r \right]_{|u=1}. \quad (7)$$

Equivalently, by substituting $(1 - \alpha_j)$ for α_0 , we have

$$\varphi(\alpha_j) = \sum_{i=0}^{\lfloor (F-1)/j \rfloor} \binom{r}{i} \alpha_j^i (1 - \alpha_j)^{r-i} \quad \text{for } rj \geq F. \quad (8)$$

Notice that for $rj < F$, the maximum degree of u in (7) is less than F . Therefore, the operator $\mathbf{d}_{<F}^u$ does not eliminate any term from the expansion and $\varphi(\alpha_j) = 1$.

Expression (8) has the form of the binomial distribution CDF; The following lemma helps us to determine its minima.

Lemma 5: The function $\varphi(\alpha_j)$ in (8) has a local minimum in all the points j where $\lfloor \frac{F-1}{j-1} \rfloor - \lfloor \frac{F-1}{j} \rfloor \geq 1$. In other words, $\varphi(\alpha_j)$ minimizes over some j^* of the form $\lceil \frac{F-1}{i} \rceil$ for some i .

Proof: For constants m and n , $f(x) = \sum_{i=0}^m \binom{n}{i} x^i (1-x)^{n-i}$ is decreasing in x . Therefore, if $j_1 < j_2$ and $\lfloor \frac{F-1}{j_1} \rfloor = \lfloor \frac{F-1}{j_2} \rfloor$, then $\alpha_{j_1} > \alpha_{j_2}$ and thus, $\varphi(\alpha_{j_1}) < \varphi(\alpha_{j_2})$. \blacksquare

Lemma 5 reduces the complexity of finding the minimum of (8) considerably, as it limits the search for optimal j , shown by j^* , to the set of per node budgets $\{\lceil \frac{F}{i} \rceil : i \in [r]\}$. Therefore, finding the optimal symmetric allocation is reduced to computing the probability of successful recovery of the original file when α_{j^*} fraction of the nodes contain j^* portion of the file and the rest of the nodes are empty.

In order to find the optimal value j^* , we derive the probability of successful decoding of a random receiver. Suppose that only d out of r of storage nodes to which a receiver has access are non-empty. In this case, the receiver can recover the file only if $d \geq i$. Therefore, the probability of successful file recovery when each non-empty storage node has $\lceil \frac{F}{i} \rceil$ portion of the file is

$$\frac{1}{\binom{r}{d}} \sum_{d=i}^r \binom{\lceil \frac{T}{i} \rceil}{d} \binom{n - \lceil \frac{T}{i} \rceil}{r-d}, \quad (9)$$

which has the form of the CDF of hyper-geometric distribution. We have to evaluate this function for all $i \in [r]$ and choose j^* such that the highest success probability is achieved. Note that given r the solution can be found in constant time since by Lemma 5 we just need to evaluate (9) r times.

A. Symmetric Allocation in Connected Random Graphs

In a practical network, a node cannot connect (via single hop) to every subset of r nodes. As a first step towards practical settings, we investigate the asymptotics of the allocation problem in large random graphs. A random graph $G(n, p)$ has n vertices, and every two vertices are connected with probability p . We direct our attention to connected random graphs since they better describe real networks. $G(n, p)$ is connected iff p is greater than a critical value $\frac{\log n}{n}$. If $p = \frac{d \log n}{n}$ for some constant d , then $G(n, \frac{d \log n}{n})$ is connected with high probability and every vertex has degree $r \asymp \log n$ [7].

Suppose that we want to store a file of size F and budget T in such a graph provided that each node could reconstruct the file by accessing its 1-hop neighbors. We are interested

in maximizing the probability that a node is successful, as the number of nodes n in the network grows. It is clear that the budget T should also grow in order to maintain a certain success probability for receivers. Otherwise, probability of successful recovery of the file will be 0. Given T , the mean number of symbols per node is T/n and therefore the mean number of symbols a node has access to is equal to $\frac{rT}{n}$. Since the file size is assumed to be constant, the most important regime to study is when $\frac{rT}{n} \asymp \mu$, where μ is a constant.

In this regime, every one of the random variables x_{s_1}, \dots, x_{s_r} , representing the number of symbols in every chosen node, is a non-negative random variable with the expectation μ/r . Standard limit theorems ([8]) imply that the random variable $\sum_{i=1}^r x_{s_i}$ will follow approximately a Poisson distribution. Consider the case $r \asymp d \log n$ and $T = \mu n/r$. For $i = 1, \dots, F$, define λ_i so that $\alpha_i = \frac{\lambda_i}{r}$ and let X_1, \dots, X_F be independent Poisson random variables such that X_i follows $\text{Poisson}(k; \lambda_i) = \lambda_i^k e^{-\lambda_i} / k!$. Then, classic approximation theorems ([9], [10]) imply that the random variables $\sum_{i=1}^r x_{s_i}$ and $\sum_{i=1}^F iX_i$ behave similarly. In fact, their difference in total variation obeys the following bound

$$\text{TV}\left(\sum_{i=1}^r x_{s_i}, \sum_{i=1}^F iX_i\right) = O\left(\frac{1}{\log n}\right).$$

Therefore, it is the case that

$$\mathbf{P}_{s \sim [n]^r} \left(\sum_{i=1}^r x_{s_i} < F \right) = \mathbf{P} \left(\sum_{i=1}^F iX_i < F \right) + O\left(\frac{1}{\log n}\right).$$

In the symmetric case, we allocate either 0 or j symbols. Hence, at most λ_0 and λ_j have non-zero values. Since in symmetric case we have $j\alpha_j = r\mu$, the previous expression becomes

$$\mathbf{P}_{s \sim [n]^r} \left(\sum_{i=1}^r x_{s_i} < F \right) = \sum_{k=0}^{\lfloor \frac{F-1}{j} \rfloor} \frac{(\mu/j)^k e^{-\mu/j}}{k!} + O\left(\frac{1}{\log n}\right).$$

Similar to the result in the previous section, since $e^{-x} \sum_{k=0}^m x^k / k!$ is a decreasing in x , in order to find the optimal j , we just need to evaluate the above expression for $j \in \{\lceil F/i \rceil : i \in [r]\}$ and the optimal value j^* is the one which maximizes the success probability.

V. SIMULATION RESULTS AND CONCLUSION

We numerically investigated the results of section IV through some simulations. Due to the complexity of the problem, finding the true optimal allocation for large n is not practical. In order to verify our results, we compare the approximate solution with optimal (found by searching all symmetric allocations) for two different examples. First, for $n = 10$ and $r = 2$, optimal symmetric allocation consists of two parts: for $T/F \in (1, 4.5)$, the file should be stored entirely and, for $T/F > 4.5$, all storage locations should store half of the file. As shown in Figure 2, approximate solution gives correct allocation for this case. For the second case, where $n = 15$ and $r = 5$, the optimal allocation is more complicated. We observe that the choice of $j/F = 1$ remains optimal until

$T/F = 4.5$. Then, for $T/F \in (4.5, 4.65)$, the optimal number of nodes to use is 9 ($= \lfloor T/j^* \rfloor$) and each of them store half of the file. Finally, we observe a transition that spreads the file maximally over all storage nodes. It is interesting that in this case our approximate solution again matches the optimal symmetric allocation. Figure 3 plots the probability of success versus normalized budget for $n = 15$ and $r = 3$




Fig. 2. Optimal symmetric allocation vs normalized budget.




Fig. 3. Probability of success vs normalized budget for $n = 15$ and $r = 3$ in the symmetric allocation where each node stores j/F fraction of the file.

In general, we observe a transition from concentration of budget over minimal number of nodes to maximal spreading of the budget over all storage nodes as the budget increases (this observation is also reported in [6]). This transition is not sharp as we observed that there are cases where the number of non-empty nodes is neither of the extremes. Also, where the transition happens is not trivial to determine and for each budget the optimal allocation should be computed using the machinery developed in this paper. Finding useful algorithms in order to find the optimal allocation in general sense and also for more realistic scenarios remains of interest.

REFERENCES

- [1] M. Naor and R. M. Roth, "Optimal file sharing in distributed networks," in *Proc. nd Annual Symposium on Foundations of Computer Science*, 1991, pp. 515–525.
- [2] A. Jiang and J. Bruck, "Memory allocation in information storage networks," in *Proc. IEEE International Symposium on Information Theory*, 2003, pp. 453–.
- [3] —, "Network file storage with graceful performance degradation," *Trans. Storage*, vol. 1, no. 2, pp. 171–189, 2005.
- [4] Z. Kong, S. A. Aly, and E. Soljanin, "Decentralized coding algorithms for distributed storage in wireless sensor networks," *CoRR*, vol. abs/0904.4057, 2009.
- [5] A. G. Dimakis, V. Prabhakaran, and K. Ramchandran, "Ubiquitous access to distributed data in large-scale sensor networks through decentralized erasure codes," in *Proc. Fourth International Symposium on Information Processing in Sensor Networks IPSN 2005*, 2005.
- [6] D. Leong, A. G. Dimakis, and T. Ho, "Distributed storage allocation problems," in *Proc. Workshop on Network Coding, Theory, and Applications NetCod '09*, 2009, pp. 86–91.

- [7] N. Alon and J. Spencer, *The Probabilistic Method—3rd edition*. John Wiley & Sons, USA, 2008.
- [8] V. Petrov, *Limit theorems of probability theory: sequences of independent random variables*. Oxford University Press, USA, 1995.
- [9] L. Le Cam, “An approximation theorem for the Poisson binomial distribution,” *Pacific J. Math*, vol. 10, no. 4, pp. 1181–1197, 1960.
- [10] A. Barbour, “Stein’s method and Poisson process convergence,” *Journal of Applied Probability*, pp. 175–184, 1988.