

Distributed Storage Allocations

Derek Leong, Alexandros G. Dimakis, and Tracey Ho

Abstract—We examine the problem of allocating a given total storage budget in a distributed storage system for maximum reliability. A source has a single data object that is to be coded and stored over a set of storage nodes; it is allowed to store any amount of coded data in each node, as long as the total amount of storage used does not exceed the given budget. A data collector subsequently attempts to recover the original data object by accessing only the data stored in a random subset of the nodes. By using an appropriate code, successful recovery can be achieved whenever the total amount of data accessed is at least the size of the original data object. The goal is to find an optimal storage allocation that maximizes the probability of successful recovery. This optimization problem is challenging in general because of its combinatorial nature, despite its simple formulation. We study several variations of the problem, assuming different allocation models and access models. The optimal allocation and the optimal symmetric allocation (in which all nonempty nodes store the same amount of data) are determined for a variety of cases. Our results indicate that the optimal allocations often have nonintuitive structure and are difficult to specify. We also show that depending on the circumstances, coding may or may not be beneficial for reliable storage.

Index Terms—Data storage systems, distributed storage, network coding, reliability, storage allocation.

I. INTRODUCTION

CONSIDER a distributed storage system comprising n storage nodes. A source has a single data object of normalized unit size that is to be coded and stored in a distributed manner over these nodes, subject to a given total storage budget T . Let x_i be the amount of coded data stored in node $i \in \{1, \dots, n\}$. Any amount of data may be stored in each node, as long as the total amount of storage used over all nodes is at most the given budget T , i.e.,

$$\sum_{i=1}^n x_i \leq T.$$

This paper is an extended version of [1], which is available online at <http://dx.doi.org/10.1109/TIT.2012.2191135>. The material in this paper was presented in part at the NetCod 2009 [2], ICC 2010 [3], and GLOBECOM 2010 [4] conferences.

The work of D. Leong and T. Ho was supported in part by Subcontract 069153 issued by BAE Systems National Security Solutions, Inc. and supported by the Defense Advanced Research Projects Agency (DARPA) and the Space and Naval Warfare System Center (SPAWARSYSCEN), San Diego, under Contract N66001-08-C-2013, by the Air Force Office of Scientific Research under Grant FA9550-10-1-0166, and by the Lee Center for Advanced Networking at the California Institute of Technology (Caltech). The work of D. Leong was supported in part by A*STAR, Singapore. The work of A. G. Dimakis was supported in part by the Center for the Mathematics of Information at Caltech, and by the National Science Foundation under NSF Grant 1055099.

D. Leong and T. Ho are with the Department of Electrical Engineering, California Institute of Technology, Pasadena, California 91125, USA (email: derekleong@caltech.edu, tho@caltech.edu).

A. G. Dimakis is with the Department of Electrical Engineering, University of Southern California, Los Angeles, California 90089, USA (email: dimakis@usc.edu).

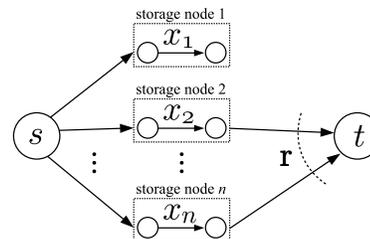


Fig. 1. Information flows in a distributed storage system. The source s has a single data object of normalized unit size that is to be coded and stored over n storage nodes. Subsequently, a data collector t attempts to recover the original data object by accessing only the data stored in a random subset r of the nodes.

This is a realistic constraint if there is limited transmission bandwidth or storage space, or if it is too costly to mirror the data object in its entirety in every node. At some time after the creation of this coded storage, a data collector attempts to recover the original data object by accessing only the data stored in a random subset r of the nodes, where the probability distribution of $r \subseteq \{1, \dots, n\}$ is specified by an assumed access model or failure model (nodes or links may fail probabilistically, for example). Fig. 1 depicts such a distributed storage system.

The reliability of this system, which we define to be the probability of successful recovery (or recovery probability in short), depends on both the storage allocation and the coding scheme. For maximum reliability, we would therefore need to find

- (i) an optimal allocation of the given budget T over the nodes, specified by the values of x_1, \dots, x_n , and
- (ii) an optimal coding scheme

that jointly maximize the probability of successful recovery. It turns out that these two problems can be decoupled by using a good coding scheme, specifically one that enables successful recovery whenever the total amount of data accessed by the data collector is at least the size of the original data object. This can be seen by considering the information flows for a network in which the source is multicasting the data object to a set of potential data collectors [5], [6]: successful recovery can be achieved by a data collector if and only if its corresponding max-flow or min-cut from the source is at least the size of the original data object. Random linear coding over a sufficiently large field would allow successful recovery with high probability when this condition is satisfied [7], [8]. Alternatively, a suitable maximum distance separable (MDS) code for the given budget and data object size would allow successful recovery with certainty when this condition is satisfied.

Therefore, assuming the use of an appropriate code, the probability of successful recovery for an allocation

(x_1, \dots, x_n) can be written as

$$\mathbb{P}[\text{successful recovery}] = \mathbb{P}\left[\sum_{i \in \mathbf{r}} x_i \geq 1\right].$$

Our goal is to find an optimal allocation that maximizes this recovery probability, subject to the given budget constraint.

Although we have assumed coded storage at the outset, coding may ultimately be unnecessary for certain allocations. For example, if the budget is spread minimally such that each nonempty node stores the data object in its entirety (i.e., $x_i \geq 1$ for all $i \in S$, and $x_i = 0$ for all $i \notin S$, where S is some subset of $\{1, \dots, n\}$), then uncoded replication would suffice since the data object can be recovered by accessing any *one* nonempty node; the data collector would not need to combine data accessed from different nodes in order to recover the data object. Thus, by solving for the optimal allocation, we will also be able to determine whether coding is beneficial for reliable storage.

We note that even though no explicit upper bound is imposed on the amount of data that can be stored in each node, it is never necessary to set $x_i > 1$ because $x_i = 1$ already allows the data object to be stored in its entirety in that node. The absence of a tighter per-node storage constraint $x_i \leq c_i < 1$ is reasonable for storage systems that handle a large number of data objects: we would expect the storage capacity of each node to be much larger than the size of a single data object, making it possible for a node to accommodate some of the data objects in their entirety. As such, it would be appropriate to apply a storage constraint for each data object via the budget T , without a separate *a priori* constraint for x_i . Furthermore, the simplifying assumption of x_i being a continuous variable is a reasonable one for large data objects: a large data object size would facilitate the creation of coded data packets with sizes (closely) matching that of a desired allocation. Incidentally, the overhead associated with random linear coding or an MDS code, which is ignored in our model, becomes proportionately negligible when the amount of coded data is large.

In spite of the simple formulation, this optimization problem poses significant challenges because of its combinatorial nature and the large space of feasible allocations. Different variations of this problem can be formulated by assuming different allocation models and access models; in this paper, we will examine three such variations that are motivated by practical storage problems in content delivery networks, delay tolerant networks, and wireless sensor networks.

A. Independent Probabilistic Access to Each Node

In the first problem formulation, we assume that the data collector accesses each of the n nodes independently with constant probability p ; in other words, each node i appears in subset \mathbf{r} independently with probability p . The resulting problem can be interpreted as that of maximizing the reliability of data storage in a system comprising n storage devices where each device fails independently with probability $1 - p$. It is not hard to show that determining the recovery probability of a *given* allocation is computationally difficult (specifically, #P-hard). The intuitive approach of spreading the budget

maximally over all nodes, i.e., setting $x_i = \frac{T}{n}$ for all i , turns out to be not necessarily optimal; in fact, the optimal allocation may not even be symmetric (we say that an allocation is *symmetric* when all nonzero x_i are equal). The following counterexample from [9] demonstrates that symmetric allocations can be suboptimal: for $(n, p, T) = (5, \frac{2}{3}, \frac{7}{3})$, the nonsymmetric allocation

$$\left(\frac{2}{3}, \frac{2}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right),$$

which achieves a recovery probability of 0.90535, performs strictly better than any symmetric allocation; the maximum recovery probability among symmetric allocations is 0.88889, which is achieved by both

$$\left(\frac{7}{6}, \frac{7}{6}, 0, 0, 0\right) \text{ and } \left(\frac{7}{12}, \frac{7}{12}, \frac{7}{12}, \frac{7}{12}, 0\right).$$

Evidently, the simple strategy of “spreading eggs evenly over more baskets” may not always improve the reliability of an allocation.

Our Contribution: We show that the intuitive symmetric allocation that spreads the budget maximally over all nodes is indeed *asymptotically* optimal in a regime of interest. Specifically, we derive an upper bound for the suboptimality of this allocation, and show that the performance gap vanishes asymptotically as the total number of storage nodes n grows, when $p > \frac{1}{T}$. This is a regime of interest because a high recovery probability is possible when $p > \frac{1}{T} \iff pT > 1$: The expected total amount of data accessed by the data collector is given by

$$\mathbb{E}\left[\sum_{i=1}^n x_i Y_i\right] = \sum_{i=1}^n x_i \mathbb{E}[Y_i] = p \sum_{i=1}^n x_i \leq pT, \quad (1)$$

where Y_i 's are independent Bernoulli(p) random variables. Therefore, the data collector would be able to access a sufficient amount of data *in expectation* for successful recovery if $pT > 1$.

We also show that the symmetric allocation that spreads the budget minimally is optimal when p is sufficiently small. In such an allocation, the data object is stored in its entirety in each nonempty node, making coding unnecessary. Additionally, we explicitly find the optimal *symmetric* allocation for a wide range of parameter values of p and T .

Related Work: This problem was introduced to us through a discussion at UC Berkeley [9]. We have since learned that variations of the problem have also been studied in several different fields.

In reliability engineering, the weighted- k -out-of- n system [10] comprises n components, each having a positive integer weight w_i and surviving independently with probability p_i ; the system is in a good state if and only if the total weight of its surviving components is at least a specified threshold k . Related work on this system and its extensions has focused on the efficient computation of the reliability of a given weight allocation (see, e.g., [11]).

In peer-to-peer networking, the allocation problem deals with the recovery of a data object from peers that are available only probabilistically. Lin *et al.* [12] compared the performance of uncoded replication vs. coded storage, restricted to

symmetric allocations, for the case where the budget is an integer.

In wireless communications, the allocation problem is studied in the context of multipath routing, in which coded data is transmitted along different paths in an unreliable network, exploiting path diversity to improve the reliability of end-to-end communications. Tsirigos and Haas [13], [14] examined the performance of symmetric allocations and noted the existence of a phase transition in the optimal symmetric allocation; approximation methods were also proposed by the authors to tackle the optimization problem, especially for the case where path failures occur with nonuniform probabilities and may be correlated. Jain *et al.* [15] evaluated the performance of symmetric allocations experimentally in a delay tolerant network setting, and presented an alternative formulation using Gaussian distributions to model partial access to nodes.

Our work generalizes these previous efforts by considering nonsymmetric allocations and noninteger budgets. We also correct some inaccurate claims about the optimal symmetric allocation in [15] and its associated technical report.

B. Access to a Random Fixed-Size Subset of Nodes

In the second problem formulation, we assume that the data collector accesses an r -subset of the n nodes selected uniformly at random from the collection of all $\binom{n}{r}$ possible r -subsets, where r is a given constant. The resulting problem can be interpreted as that of maximizing the recovery probability in a networked storage system of n nodes where the end user is able or allowed to contact up to r nodes randomly. We can treat this access model as an approximation to the preceding independent probabilistic access model by picking $r \approx np$. Finding the optimal allocation in this case is still challenging. As in the first problem formulation, it is not hard to show that determining the recovery probability of a given allocation is computationally difficult (specifically, #P-complete).

The problem appears nontrivial even if we restrict the optimization to only *symmetric* allocations. Numerically, we observe that given n and r , either a minimal or a maximal spreading of the budget is optimal among symmetric allocations for most, if not all, choices of T . One example of an exception is $(n, r, T) = (14, 5, \frac{8}{3})$ for which it is optimal to have 8 nonempty nodes in the symmetric allocation, instead of the extremes 2 or 13; another example is $(n, r, T) = (16, 4, \frac{7}{2})$ for which it is optimal to have 7 nonempty nodes in the symmetric allocation, instead of the extremes 3 or 14. Furthermore, the number of nonempty nodes in the optimal symmetric allocation is not necessarily a nondecreasing function of the budget T ; for instance, given $(n, r) = (20, 4)$, it is optimal to have $(4, 18, 14, 19, 20)$ nonempty nodes in the symmetric allocation for $T = (4.25, 4.5, 4.67, 4.75, 5)$, respectively.

Our Contribution: We show that the allocation $(\frac{1}{r}, \dots, \frac{1}{r})$ is optimal in the *high recovery probability regime*. Specifically, we demonstrate that this allocation, which has a recovery probability of exactly 1, minimizes the budget T necessary for achieving any recovery probability exceeding a specified threshold $1 - \epsilon$. Although ϵ depends on n and r in a compli-

cated way, we can conclude that for any r , this allocation is optimal if the recovery probability is to exceed $1 - \frac{1}{n}$.

We also make the following conjecture about the optimal allocation, based on our numerical observations:

Conjecture. A *symmetric* optimal allocation always exists for any n , r , and T .

Related Work: Sardari *et al.* [16] presented a method of *approximating* an optimal solution to this problem by considering a data collector that accesses r random nodes with replacement. More recently, Alon *et al.* [17] showed that this problem is related to an old conjecture by Erdős on the maximum number of edges in a uniform hypergraph [18].

C. Probabilistic Symmetric Allocations

In the third problem formulation, we assume a *probabilistic* allocation model in which the source selects a random allocation from a distribution of allocations, with the constraint that the *expected* total amount of storage used in an allocation is at most the given budget T . We specifically consider the case where each of the n nodes is selected by the source independently with constant probability $\min(\frac{\ell T}{n}, 1)$ to store a constant $\frac{1}{\ell}$ amount of data, thus creating a probabilistic *symmetric* allocation of the budget. The data collector subsequently accesses an r -subset of the n nodes selected uniformly at random from the collection of all $\binom{n}{r}$ possible r -subsets, where r is a given constant. The goal is to find an optimal allocation, specified by the value of parameter ℓ , that maximizes the recovery probability. This model was conceived as a simplification of the preceding fixed-size subset access model which assumes a deterministic allocation of the budget.

Our Contribution: We show that the choice of $\ell = r$, which corresponds to a maximal spreading of the budget, is optimal when the given budget T is sufficiently large, or equivalently, when a sufficiently high recovery probability (specifically, $\frac{3}{4}$ or higher) is achievable. We believe this is a reasonable operating regime for applications that require good reliability.

D. Other Related Work

Apart from the work done on the preceding problems, a variety of storage allocation problems have also been studied in a *nonprobabilistic* setting. For instance, the objective adopted in [19] and [20] is to minimize the total storage budget required to satisfy a given set of deterministic recovery requirements in a network. Incidentally, the use of network coding makes it easier to deal with the total cost of content delivery, which covers the initial dissemination, storage, and eventual fetching of data objects; this cost-minimization problem is considered in [6] and [21], subject to various deterministic constraints involving, for example, load balancing or fetching distance.

We note that in most of the literature involving reliable distributed storage, either the data object is assumed to be replicated in its entirety (see, e.g., [22]), or, if coding is used, every node is assumed to store the same amount of coded data (see, e.g., [23]–[27]). Allocations of a storage budget with nodes possibly storing different amounts of data are not usually considered.

TABLE I
NOTATION

Symbol	Definition
n	total number of storage nodes, $n \geq 2$
x_i	amount of data stored in storage node i , $x_i \geq 0$, where $i \in \{1, \dots, n\}$
T	total storage budget, $1 \leq T \leq n$
\mathbf{r}	subset of nodes accessed, $\mathbf{r} \subseteq \{1, \dots, n\}$
p	access probability (Section II), $0 < p < 1$
r	number of nodes accessed (Section III), $1 \leq r \leq n$
$\frac{1}{\ell}$	amount of data stored in each nonempty node (Section IV), $\ell > 0$
$\mathcal{B}(n, p)$	binomial random variable with n trials and success probability p
$\mathbb{1}[G]$	indicator function; $\mathbb{1}[G] = 1$ if statement G is true, and 0 otherwise
\mathbb{Z}_0^+	the set of nonnegative integers, i.e., $\mathbb{Z}^+ \cup \{0\}$

In the following three sections, we define each problem formally and state our main results. Proofs of theorems are deferred to the appendix. Table I summarizes the notation used throughout this paper.

II. INDEPENDENT PROBABILISTIC ACCESS TO EACH NODE

In the first variation of the storage allocation problem, we consider a data collector that accesses each of the n nodes independently with probability p ; successful recovery occurs if and only if the total amount of data stored in the accessed nodes is at least 1. We seek an optimal allocation (x_1, \dots, x_n) of the budget T that maximizes the probability of successful recovery, for a given choice of n , p , and T . This optimization problem can be expressed as follows:

$\Pi_1(n, p, T) :$

$$\text{maximize}_{x_1, \dots, x_n} \sum_{\mathbf{r} \subseteq \{1, \dots, n\}} p^{|\mathbf{r}|} (1-p)^{n-|\mathbf{r}|} \cdot \mathbb{1} \left[\sum_{i \in \mathbf{r}} x_i \geq 1 \right] \quad (2)$$

subject to

$$\sum_{i=1}^n x_i \leq T \quad (3)$$

$$x_i \geq 0 \quad \forall i \in \{1, \dots, n\}. \quad (4)$$

The objective function (2) is just the recovery probability, expressed as the sum of the probabilities corresponding to the subsets \mathbf{r} that allow successful recovery. An equivalent expression for (2) is

$$\mathbb{P} \left[\sum_{i=1}^n x_i Y_i \geq 1 \right],$$

where Y_i 's are independent Bernoulli(p) random variables. Inequality (3) expresses the budget constraint, and inequality (4) ensures that a nonnegative amount of data is stored in each node. For the trivial budget $T = 1$, the allocation $(1, 0, \dots, 0)$ is optimal; for $T = n$, the allocation $(1, \dots, 1)$ is optimal.

TABLE II
OPTIMAL ALLOCATIONS FOR NUMBER OF NODES $n = 2, 3, 4$

n	Budget T	Optimal allocation	Condition on access probability p (if any)
2	$1 \leq T < 2$	(1, 0)	
	$1 \leq T < \frac{3}{2}$	(1, 0, 0)	
3	$\frac{3}{2} \leq T < 2$	(1, 0, 0) $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$	if $p \leq \frac{1}{2}$ if $p \geq \frac{1}{2}$
	$2 \leq T < 3$	(1, 1, 0)	
	$1 \leq T < \frac{4}{3}$	(1, 0, 0, 0)	
4	$\frac{4}{3} \leq T < \frac{3}{2}$	(1, 0, 0, 0) $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, \frac{1}{3})$	if $p \leq \frac{1+\sqrt{13}}{6} \approx 0.768$ if $p \geq \frac{1+\sqrt{13}}{6} \approx 0.768$
	$\frac{3}{2} \leq T < 2$	(1, 0, 0, 0) $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, 0)$	if $p \leq \frac{1}{2}$ if $p \geq \frac{1}{2}$
	$2 \leq T < \frac{5}{2}$	(1, 1, 0, 0) $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2})$	if $p \leq \frac{2}{3}$ if $p \geq \frac{2}{3}$
	$\frac{5}{2} \leq T < 3$	(1, 1, 0, 0) $(1, \frac{1}{2}, \frac{1}{2}, \frac{1}{2})$	if $p \leq \frac{1}{2}$ if $p \geq \frac{1}{2}$
	$3 \leq T < 4$	(1, 1, 1, 0)	

Incidentally, computing the recovery probability of a *given* allocation turns out to be #P-hard:

Proposition 1. *Computing the recovery probability*

$$\sum_{\mathbf{r} \subseteq \{1, \dots, n\}} p^{|\mathbf{r}|} (1-p)^{n-|\mathbf{r}|} \cdot \mathbb{1} \left[\sum_{i \in \mathbf{r}} x_i \geq 1 \right]$$

for a given allocation (x_1, \dots, x_n) and choice of p is #P-hard.

Table II lists the optimal allocations for $n = 2, 3, 4$, covering all parameter values of $p \in (0, 1)$ and $T \in [1, n]$. These solutions are obtained by minimizing T for each possible value of the objective function (2). We observe that

- (i) for any T , the symmetric allocation $(1, \dots, 1, 0, \dots, 0)$, which corresponds to a minimal spreading of the budget (uncoded replication), appears to be optimal when p is sufficiently small, and
- (ii) the optimal *symmetric* allocation appears to perform well despite being suboptimal in some cases, e.g., when $(n, T) = (4, \frac{5}{2})$ and $p > \frac{1}{2}$.

We will proceed to show that observation (i) is indeed true in Section II-B; the opposite approach of spreading the budget maximally over all nodes turns out to be *asymptotically* optimal when p is sufficiently large, as will be demonstrated in Section II-A. Motivated by observation (ii), we examine the optimization problem restricted to symmetric allocations in Section II-C.

For brevity, let $\bar{x}(n, T, m)$ denote the *symmetric* allocation for n nodes that uses a total storage of T and contains exactly

$m \in \{1, 2, \dots, n\}$ nonempty nodes:

$$\bar{\mathbf{x}}(n, T, m) \triangleq \left(\underbrace{\frac{T}{m}, \dots, \frac{T}{m}}_{m \text{ entries}}, \underbrace{0, \dots, 0}_{(n-m) \text{ entries}} \right).$$

Since successful recovery for the symmetric allocation $\bar{\mathbf{x}}(n, T, m)$ occurs if and only if at least $\lceil 1 / (\frac{T}{m}) \rceil = \lceil \frac{m}{T} \rceil$ out of the m nonempty nodes are accessed, the corresponding probability of successful recovery can be written as

$$P_S(p, T, m) \triangleq \mathbb{P} \left[\mathcal{B}(m, p) \geq \left\lceil \frac{m}{T} \right\rceil \right].$$

A. Asymptotic Optimality of Maximal Spreading

The recovery probability of the symmetric allocation $\bar{\mathbf{x}}(n, T, m=n)$, which corresponds to a maximal spreading of the budget over all nodes, is given by

$$P_S(p, T, m=n) = \mathbb{P} \left[\mathcal{B}(n, p) \geq \left\lceil \frac{n}{T} \right\rceil \right]. \quad (5)$$

To establish the optimality of this allocation, we compare (5) to an upper bound for the recovery probability of an optimal allocation. Such a bound can be derived by conditioning on the number of accessed nodes:

Lemma 1. *The probability of successful recovery for an optimal allocation is at most*

$$\sum_{r=0}^n \min \left(\frac{rT}{n}, 1 \right) \mathbb{P} [\mathcal{B}(n, p) = r]. \quad (6)$$

The suboptimality of $\bar{\mathbf{x}}(n, T, m=n)$ is therefore bounded by the difference between (5) and (6), as given by the following theorem; when $p > \frac{1}{T}$, this allocation becomes asymptotically optimal since its suboptimality gap vanishes as n goes to infinity:

Theorem 1. *The gap between the probabilities of successful recovery for an optimal allocation and for the symmetric allocation $\bar{\mathbf{x}}(n, T, m=n)$, which corresponds to a maximal spreading of the budget over all nodes, is at most*

$$pT \mathbb{P} \left[\mathcal{B}(n-1, p) \leq \left\lceil \frac{n}{T} \right\rceil - 2 \right].$$

If p and T are fixed such that $p > \frac{1}{T}$, then this gap approaches zero as n goes to infinity.

We note that the regime $p > \frac{1}{T}$ is particularly interesting because it corresponds to the regime of high recovery probability; the recovery probability would be bounded away from 1 if $p < \frac{1}{T} \iff pT < 1$ instead. This follows from the application of Markov's inequality to the random variable W denoting the total amount of data accessed by the data collector, which produces

$$\mathbb{P} [W \geq 1] \leq \mathbb{E} [W].$$

Since $\mathbb{P} [W \geq 1]$ is just the probability of successful recovery, and $\mathbb{E} [W] \leq pT$ according to (1), we have

$$\mathbb{P} [\text{successful recovery}] \leq pT.$$

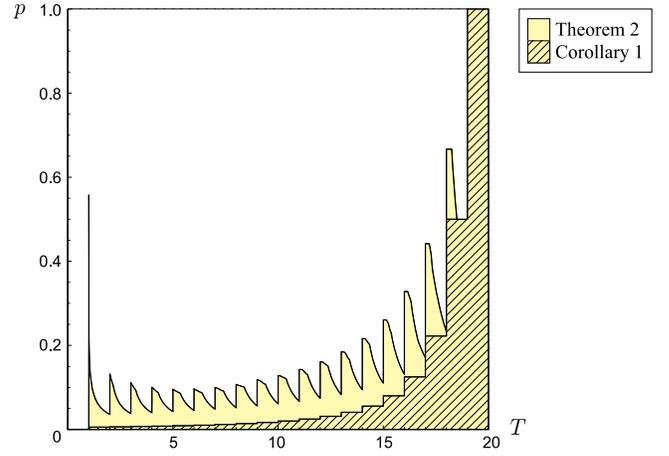


Fig. 2. Plot of access probability p against budget T , showing regions of (T, p) over which the sufficient conditions of the theorems are satisfied, for $n = 20$. Minimal spreading (uncoded replication) is optimal among all allocations in the colored regions.

B. Optimality of Minimal Spreading (Uncoded Replication)

The recovery probability of the symmetric allocation $\bar{\mathbf{x}}(n, T, m=\lfloor T \rfloor)$, which corresponds to a minimal spreading of the budget, is given by

$$P_S(p, T, m=\lfloor T \rfloor) = \mathbb{P} [\mathcal{B}(\lfloor T \rfloor, p) \geq 1] = 1 - (1-p)^{\lfloor T \rfloor}. \quad (7)$$

Recall that coding is unnecessary in such an allocation since the data object is stored in its entirety in each nonempty node. A sufficient condition for the optimality of this allocation can be found by comparing (7) to an upper bound for the recovery probabilities of all other allocations. Our approach is to classify each allocation according to the number of individual nodes that store at least a unit amount of data. We then find a bound for allocations containing exactly 0 such nodes, another bound for allocations containing exactly 1 such node, and so on. The subsequent comparisons of (7) to each of these bounds result in the following theorem:

Theorem 2. *If $1 < T < n$ and*

$$1 - (1-p)^{\lfloor T \rfloor - n} + (n-\ell) \left(\frac{p}{1-p} \right) + \sum_{r=2}^{\lfloor \frac{n-\ell}{T-\ell} \rfloor - 1} \left(1 - \frac{T-\ell}{n-\ell} \cdot r \right) \binom{n-\ell}{r} \left(\frac{p}{1-p} \right)^r \geq 0 \quad (8)$$

for all $\ell \in \{0, 1, \dots, \lfloor T \rfloor - 1\}$, then $\bar{\mathbf{x}}(n, T, m=\lfloor T \rfloor)$, which corresponds to a minimal spreading of the budget (uncoded replication), is an optimal allocation.

The following corollary shows that this allocation is indeed optimal for sufficiently small p :

Corollary 1. *If $1 < T < n$ and $p \leq \frac{2}{(n-\lfloor T \rfloor)^2}$, then $\bar{\mathbf{x}}(n, T, m=\lfloor T \rfloor)$ is an optimal allocation.*

Fig. 2 illustrates these results in the form of a region plot for an instance of n .

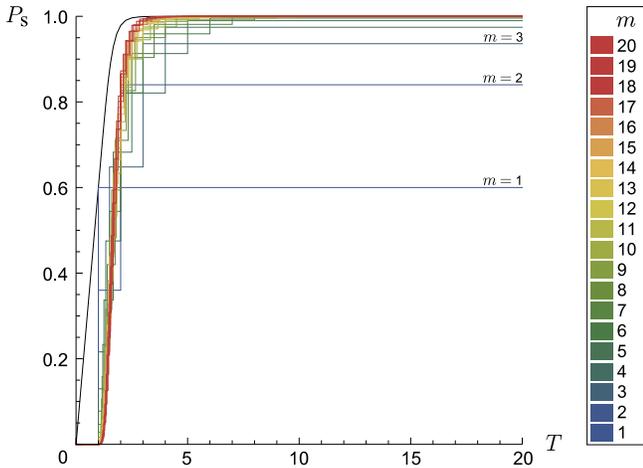


Fig. 3. Plot of recovery probability P_S against budget T for each symmetric allocation $\bar{x}(n, T, m)$, for $(n, p) = (20, \frac{2}{5})$. Parameter m denotes the number of nonempty nodes in the symmetric allocation. The black curve gives an upper bound for the recovery probability of an optimal allocation, as derived in Lemma 1.

C. Optimal Symmetric Allocation

The optimization problem appears nontrivial even if we were to consider only *symmetric* allocations. Fig. 3, which compares the performance of different symmetric allocations over different budgets for an instance of (n, p) , demonstrates that the value of m corresponding to the optimal symmetric allocation can change drastically as the budget T varies.

Fortunately, we can eliminate many candidates for the optimal value of m by making the following observation: Recall that the recovery probability of the symmetric allocation $\bar{x}(n, T, m)$ is given by $P_S(p, T, m) \triangleq \mathbb{P}[\mathcal{B}(m, p) \geq \lceil \frac{m}{T} \rceil]$. For fixed n, p , and T , we have

$$\lceil \frac{m}{T} \rceil = k \quad \text{when } m \in ((k-1)T, kT],$$

for $k = 1, 2, \dots, \lfloor \frac{n}{T} \rfloor$, and finally,

$$\lceil \frac{m}{T} \rceil = \lfloor \frac{n}{T} \rfloor + 1 \quad \text{when } m \in \left(\lfloor \frac{n}{T} \rfloor T, n \right].$$

Since $\mathbb{P}[\mathcal{B}(m, p) \geq k]$ is nondecreasing in m for constant p and k , it follows that $P_S(p, T, m)$ is maximized within each of these intervals of m when we pick m to be the largest integer in the corresponding interval. Thus, given n, p , and T , we can find an optimal m^* that maximizes $P_S(p, T, m)$ over all m from among $\lceil \frac{n}{T} \rceil$ candidates:

$$\left\{ \lfloor T \rfloor, \lfloor 2T \rfloor, \dots, \left\lfloor \lfloor \frac{n}{T} \rfloor T \right\rfloor, n \right\}. \quad (9)$$

For $m = \lfloor kT \rfloor$, where $k \in \mathbb{Z}^+$, the corresponding probability of successful recovery is given by

$$P_S(p, T, m = \lfloor kT \rfloor) = \mathbb{P}[\mathcal{B}(\lfloor kT \rfloor, p) \geq k].$$

The difference between the probabilities of successful recovery for consecutive values of $k \in \mathbb{Z}^+$ can be written as

$$\begin{aligned} \Delta(p, T, k) &\triangleq P_S(p, T, m = \lfloor (k+1)T \rfloor) - P_S(p, T, m = \lfloor kT \rfloor) \\ &= \mathbb{P}[\mathcal{B}(\lfloor (k+1)T \rfloor, p) \geq k+1] - \mathbb{P}[\mathcal{B}(\lfloor kT \rfloor, p) \geq k] \end{aligned}$$

$$\begin{aligned} &\min(\alpha_{k,T-1}, k) \\ &= \sum_{i=1}^{\min(\alpha_{k,T-1}, k)} \mathbb{P}[\mathcal{B}(\lfloor kT \rfloor, p) = k-i] \cdot \mathbb{P}[\mathcal{B}(\alpha_{k,T}, p) \geq i+1] \\ &\quad - \mathbb{P}[\mathcal{B}(\lfloor kT \rfloor, p) = k] \cdot \mathbb{P}[\mathcal{B}(\alpha_{k,T}, p) = 0], \end{aligned}$$

where $\alpha_{k,T} \triangleq \lfloor (k+1)T \rfloor - \lfloor kT \rfloor$. The above expression is obtained by comparing the branches of the probability tree for $\lfloor kT \rfloor$ vs. $\lfloor (k+1)T \rfloor$ independent Bernoulli(p) trials: the first term describes unsuccessful events (“ $\mathcal{B}(\lfloor kT \rfloor, p) < k$ ”) becoming successful (“ $\mathcal{B}(\lfloor (k+1)T \rfloor, p) \geq k+1$ ”) after the additional $\alpha_{k,T}$ trials, while the second term describes successful events (“ $\mathcal{B}(\lfloor kT \rfloor, p) \geq k$ ”) becoming unsuccessful (“ $\mathcal{B}(\lfloor (k+1)T \rfloor, p) < k+1$ ”) after the additional $\alpha_{k,T}$ trials. After further simplification, we arrive at

$$\begin{aligned} \Delta(p, T, k) &= p^k(1-p)^{\lfloor (k+1)T \rfloor - k} \\ &\quad \left\{ \sum_{i=1}^{\min(\alpha_{k,T-1}, k)} \sum_{j=i+1}^{\alpha_{k,T}} \binom{\lfloor kT \rfloor}{k-i} \binom{\alpha_{k,T}}{j} \left(\frac{p}{1-p}\right)^{-i+j} - \binom{\lfloor kT \rfloor}{k} \right\}. \end{aligned} \quad (10)$$

The following theorem essentially provides a sufficient condition on p and T for $\Delta(p, T, k) \geq 0$ for any $k \in \mathbb{Z}^+$, thereby eliminating all but the two largest candidate values for m^* in (9), i.e., $m = \lfloor \lfloor \frac{n}{T} \rfloor T \rfloor$ and $m = n$, which correspond to a maximal spreading of the budget over (almost) all nodes (they are identical when $\frac{n}{T} \in \mathbb{Z}^+$, i.e., $T = n, \frac{n}{2}, \frac{n}{3}, \dots$):

Theorem 3. *If*

$$(1-p)^{\lfloor T \rfloor} + 2\lfloor T \rfloor p(1-p)^{\lfloor T \rfloor - 1} - 1 \leq 0, \quad (11)$$

then either $\bar{x}(n, T, m = \lfloor \lfloor \frac{n}{T} \rfloor T \rfloor)$ or $\bar{x}(n, T, m = n)$, which correspond to a maximal spreading of the budget, is an optimal symmetric allocation.

The following corollary restates Theorem 3 in a slightly weaker but more convenient form:

Corollary 2. *If $p \geq \frac{4}{3\lfloor T \rfloor}$, then either $\bar{x}(n, T, m = \lfloor \lfloor \frac{n}{T} \rfloor T \rfloor)$ or $\bar{x}(n, T, m = n)$ is an optimal symmetric allocation.*

The following lemma mirrors Theorem 3 by providing a sufficient condition on p and T for $\Delta(p, T, k) \leq 0$ for any $k \in \mathbb{Z}^+$, thereby eliminating all but the smallest candidate value for m^* in (9), i.e., $m = \lfloor T \rfloor$, which corresponds to a minimal spreading of the budget (uncoded replication):

Lemma 2. *If $T > 1$, and either*

$$T = \frac{1}{p} \in \mathbb{Z}^+ \quad (12)$$

or

$$T < \frac{1}{p} \quad \text{and} \quad p(1-p)^{\lfloor T \rfloor - 1} \leq \frac{1}{T} \left(1 - \frac{1}{T}\right)^{\lfloor T \rfloor - 1}, \quad (13)$$

then $\bar{x}(n, T, m = \lfloor T \rfloor)$ is an optimal symmetric allocation.

The following lemma restates Lemma 2 in a slightly weaker but more convenient form:

Lemma 3. *If $p \leq \frac{2}{\lfloor T \rfloor} - \frac{1}{T}$, then $\bar{x}(n, T, m = \lfloor T \rfloor)$ is an optimal symmetric allocation.*

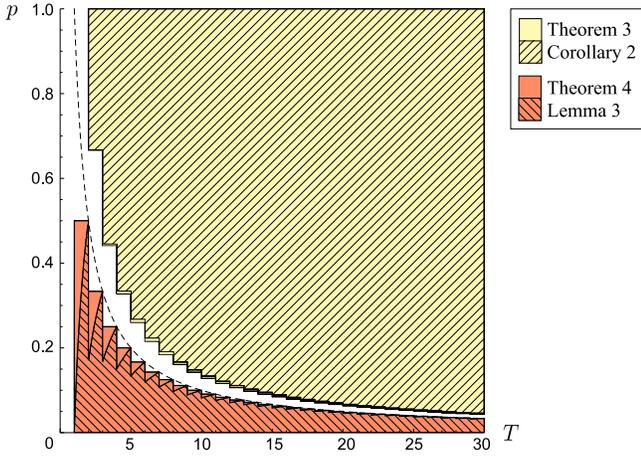


Fig. 4. Plot of access probability p against budget T , showing regions of (T, p) over which the sufficient conditions of the theorems are satisfied. The black dashed curve marks the points satisfying $p = \frac{1}{T}$. Maximal spreading is optimal among symmetric allocations in the colored regions above the curve, while minimal spreading (uncoded replication) is optimal among symmetric allocations in the colored regions below the curve.

The following theorem expands the region covered by Lemma 3 by showing that $\bar{x}(n, T, m=\lfloor T \rfloor)$ remains optimal between the “peaks” in Fig. 4:

Theorem 4. *If $p \leq \frac{1}{\lfloor T \rfloor}$, then $\bar{x}(n, T, m=\lfloor T \rfloor)$, which corresponds to a minimal spreading of the budget (uncoded replication), is an optimal symmetric allocation.*

Fig. 4 illustrates these results in the form of a region plot. The theorems cover all choices of p and T except for the gap around $p = \frac{1}{T}$, which diminishes with increasing T . Both minimal and maximal spreading of the budget may be suboptimal among symmetric allocations in this gap on either side of the curve $p = \frac{1}{T}$: for example, when $(n, p, T) = (10, \frac{9}{25}, \frac{5}{2})$, for which $p < \frac{1}{T}$, the optimal symmetric allocation is $\bar{x}(n, T, m=\lfloor 2T \rfloor)$; when $(n, p, T) = (10, \frac{3}{5}, \frac{12}{5})$, for which $p > \frac{1}{T}$, the optimal symmetric allocation is $\bar{x}(n, T, m=\lfloor 3T \rfloor)$. In general, for any budget $T \geq 2$, the optimal symmetric allocation changes from minimal spreading to maximal spreading eventually, as the access probability p increases. This transition, which is not necessarily sharp, appears to occur at around $p = \frac{1}{T}$. Interestingly, when $p = \frac{1}{T}$ exactly, we observe numerically that $\bar{x}(n, T, m=\lfloor T \rfloor)$ is the optimal symmetric allocation for *most* values of T ; the optimal symmetric allocation changes continually over the intervals

$$1.5 \leq T < 2 \quad \text{and} \quad 2.5 \leq T \leq 2.8911,$$

while $\bar{x}(n, T, m=\lfloor 2T \rfloor)$ is optimal for $3.5 \leq T \leq 3.5694$. These findings suggest that it may be difficult to specify an optimal symmetric allocation for values of p and T in the gap; we can, however, restrict our search for an optimal symmetric allocation to the $\lfloor \frac{n}{T} \rfloor$ candidates given by (9).

III. ACCESS TO A RANDOM FIXED-SIZE SUBSET OF NODES

In the second variation of the storage allocation problem, we consider a data collector that accesses an r -subset of the n nodes selected uniformly at random from the collection of all

$\binom{n}{r}$ possible r -subsets, where r is a given constant; successful recovery occurs if and only if the total amount of data stored in the accessed nodes is at least 1. We seek an optimal allocation (x_1, \dots, x_n) of the budget T that maximizes the probability of successful recovery, for a given choice of n , r , and T . This optimization problem can be expressed as follows:

$$\begin{aligned} \Pi_2(n, r, T) : \\ \text{maximize} \quad & P_S \\ \text{subject to} \quad & x_1, \dots, x_n, P_S \end{aligned} \quad (14)$$

$$\sum_{\substack{\mathbf{r} \subseteq \{1, \dots, n\} \\ |\mathbf{r}|=r}} \frac{1}{\binom{n}{r}} \cdot \mathbb{1} \left[\sum_{i \in \mathbf{r}} x_i \geq 1 \right] \geq P_S \quad (15)$$

$$\sum_{i=1}^n x_i \leq T \quad (16)$$

$$x_i \geq 0 \quad \forall i \in \{1, \dots, n\}. \quad (17)$$

The left-hand side of inequality (15) is just the recovery probability, expressed as the sum of the probabilities corresponding to the r -subsets \mathbf{r} that allow successful recovery. The objective function (14) is therefore equal to the recovery probability since P_S is maximized when (15) holds with equality. Inequality (16) expresses the budget constraint, and inequality (17) ensures that a nonnegative amount of data is stored in each node. For the trivial budget $T = 1$, the allocation $(1, 0, \dots, 0)$ is optimal; for $T \geq \frac{n}{r}$, the allocation $(\frac{1}{r}, \dots, \frac{1}{r})$, which has the maximal recovery probability of 1, is optimal. Incidentally, computing the recovery probability of a given allocation turns out to be #P-complete:

Proposition 2. *Computing the recovery probability*

$$\sum_{\substack{\mathbf{r} \subseteq \{1, \dots, n\} \\ |\mathbf{r}|=r}} \frac{1}{\binom{n}{r}} \cdot \mathbb{1} \left[\sum_{i \in \mathbf{r}} x_i \geq 1 \right]$$

for a given allocation (x_1, \dots, x_n) and choice of r is #P-complete.

An alternate way of formulating this problem is to minimize the budget T required to achieve a desired recovery probability P_S :

$$\begin{aligned} \Pi'_2(n, r, P_S) : \\ \text{minimize} \quad & T \\ \text{subject to} \quad & x_1, \dots, x_n, T \end{aligned} \quad (17)$$

subject to the three constraints (15)–(17) of $\Pi_2(n, r, T)$.

Fig. 5 shows how the optimal recovery probability $\max P_S$ varies with the budget T , for two instances of (n, r) . These plots are obtained by solving $\Pi'_2(n, r, P_S)$ for each possible value of P_S . We observe that when the budget T drops below $\frac{n}{r}$, the optimal recovery probability $\max P_S$ is reduced by a significant margin below 1. In other words, if the desired recovery probability P_S in $\Pi'_2(n, r, P_S)$ is sufficiently high, then the optimal allocation is $(\frac{1}{r}, \dots, \frac{1}{r})$, which requires a budget of $T = \frac{n}{r}$. In Section III-A, we examine the optimality of this allocation for the high recovery probability regime.

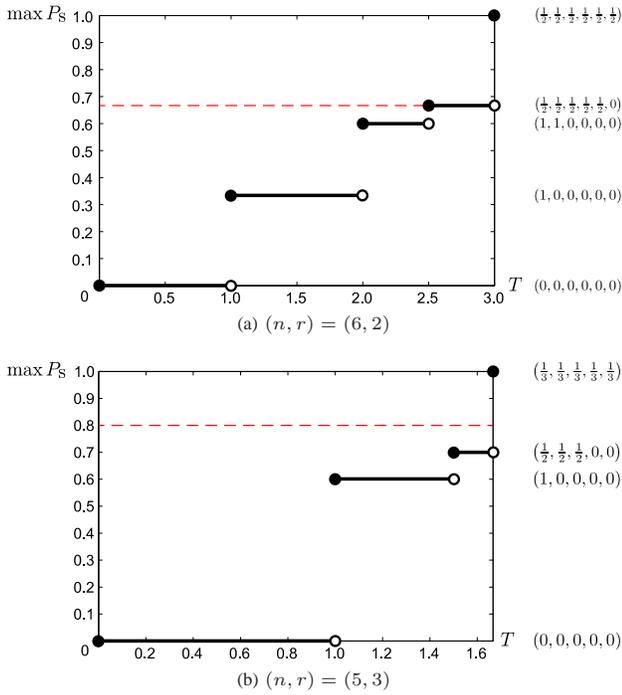


Fig. 5. Plot of the optimal recovery probability $\max P_S$ against budget T , for (a) $(n, r) = (6, 2)$ and (b) $(n, r) = (5, 3)$. The optimal allocation corresponding to each value of $\max P_S$ is given on the right-hand side of the plot. In (a), the red dashed line marks the threshold on P_S derived in Theorem 5; the allocation $(\frac{1}{2}, \dots, \frac{1}{2})$ is optimal for $\Pi'_2(n, r, P_S)$ if and only if the desired recovery probability P_S exceeds this threshold. In (b), the red dashed line marks the threshold on P_S derived in Theorem 6; the allocation $(\frac{1}{r}, \dots, \frac{1}{r})$ is optimal for $\Pi'_2(n, r, P_S)$ if P_S exceeds this threshold.

A. Regime of High Recovery Probability

Consider the optimization problem $\Pi'_2(n, r, P_S)$. We will demonstrate that the allocation $(\frac{1}{r}, \dots, \frac{1}{r})$ is optimal when the desired recovery probability P_S exceeds a specified threshold expressed in terms of n and r . Our results follow from the observation that successful recovery for certain combinations of r -subsets of nodes can impose a lower bound on the required budget T . For example, given $(n, r) = (4, 2)$, if successful recovery is to occur for $\{1, 2\}$ and $\{3, 4\}$, possibly among other r -subsets of nodes, then we have

$$\sum_{i \in \{1, 2\}} x_i \geq 1 \quad \text{and} \quad \sum_{i \in \{3, 4\}} x_i \geq 1,$$

which would imply that the minimum budget T must be at least 2, since

$$T \geq \sum_{i=1}^4 x_i = \sum_{i \in \{1, 2\}} x_i + \sum_{i \in \{3, 4\}} x_i \geq 2.$$

This observation is generalized by the following lemma:

Lemma 4. Consider a set $S \subseteq \{1, \dots, n\}$, and c subsets of S given by $\mathbf{r}_j \subseteq S$, $j = 1, \dots, c$. If

$$\sum_{i \in \mathbf{r}_j} x_i \geq 1 \quad \forall j \in \{1, \dots, c\}, \quad (18)$$

and each element in S appears exactly $b > 0$ times among the

c subsets, i.e.,

$$\sum_{j=1}^c \mathbb{1}[i \in \mathbf{r}_j] = b \quad \forall i \in S, \quad (19)$$

then

$$\sum_{i \in S} x_i \geq \frac{c}{b}.$$

We begin with the special case of probability-1 recovery, i.e., $P_S = 1$. The resulting optimization problem is just a linear program with all $\binom{n}{r}$ possible r -subset constraints.

Lemma 5. If $P_S = 1$, then $(\frac{1}{r}, \dots, \frac{1}{r})$ is an optimal allocation.

When the desired recovery probability P_S is less than 1, we can afford to drop *some* of the r -subset constraints from this linear program (recall that the recovery probability of an allocation is just the fraction of these $\binom{n}{r}$ constraints that are satisfied). Our task is to determine how many such constraints can be dropped before the lower bound for T obtained with the help of Lemma 4 falls below $\frac{n}{r}$, in which case the allocation $(\frac{1}{r}, \dots, \frac{1}{r})$ may no longer be optimal. We do this by constructing collections of r -subset constraints that yield the required lower bound of $\frac{n}{r}$ for T , and counting how many r -subset constraints need to be removed from the linear program before no such collection remains. Our answer depends on the divisibility of n by r .

When n is a multiple of r , we are able to state a necessary and sufficient condition on P_S for the allocation to be optimal:

Theorem 5. If n is a multiple of r , then $(\frac{1}{r}, \dots, \frac{1}{r})$ is an optimal allocation if and only if

$$P_S > 1 - \frac{r}{n}.$$

When n is *not* a multiple of r , we are only able to state a sufficient condition on P_S for the allocation to be optimal:

Theorem 6. If n is not a multiple of r , then $(\frac{1}{r}, \dots, \frac{1}{r})$ is an optimal allocation if

$$P_S > 1 - \frac{\gcd(r, r')}{\alpha \gcd(r, r') + r'},$$

where α and r' are uniquely defined integers satisfying

$$n = \alpha r + r', \quad \alpha \in \mathbb{Z}_0^+, \quad r' \in \{r+1, \dots, 2r-1\}.$$

However, if n is a multiple of $(n-r)$, then this sufficient condition becomes necessary too:

Corollary 3. If n is a multiple of $(n-r)$, then $(\frac{1}{r}, \dots, \frac{1}{r})$ is an optimal allocation if and only if

$$P_S > \frac{r}{n}.$$

Note that Corollary 3 allows us to solve $\Pi_2(n, r, T)$ completely when n is a multiple of $(n-r)$: for any $T \in [1, \frac{n}{r})$, the allocation $(1, 0, \dots, 0)$ is optimal since it has a recovery probability of $\frac{\binom{n-1}{r-1}}{\binom{n}{r}} = \frac{r}{n}$, i.e., exactly the threshold in Corollary 3; higher recovery probabilities are not achievable unless $T \geq \frac{n}{r}$.

Fig. 6 illustrates these results for an instance of n .

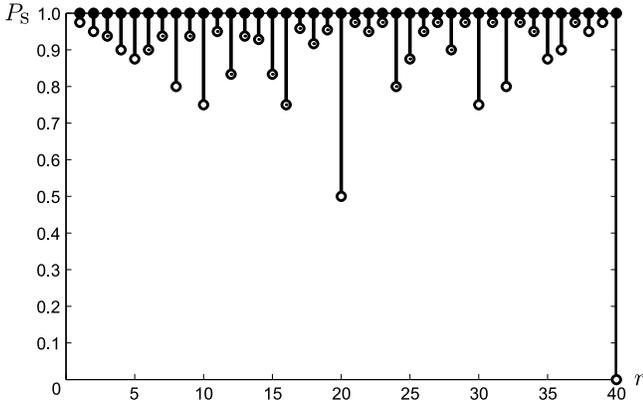


Fig. 6. Plot of the desired recovery probability P_S against the number of nodes accessed r , showing intervals of P_S over which the allocation $(\frac{1}{r}, \dots, \frac{1}{r})$ is optimal for $\Pi'_2(n, r, P_S)$, for $n = 40$. A dotted circle marker denotes an endpoint that may not be tight, i.e., we have not demonstrated that the allocation is suboptimal everywhere outside the interval.

By combining the proof techniques of Lemma 1 and Theorems 2, 5, and 6, we can derive the improved upper bound P_S^{UB} , given by (20) at the bottom of the page, for the recovery probability of an optimal allocation in the independent probabilistic access model of Section II (cf. Lemma 1). Variables α and r' are uniquely defined integers satisfying

$$n - \ell = \alpha r + r', \quad \alpha \in \mathbb{Z}_0^+, \quad r' \in \{r, \dots, 2r - 1\}.$$

Parameter ℓ denotes the number of individual nodes that store at least a unit amount of data. At least ℓ amount of data is stored in these *complete* nodes, leaving the remaining budget of at most $T - \ell$ to be allocated over the remaining $n - \ell$ *incomplete* nodes. Term (i) gives the probability of successful recovery from accessing at least one complete node, while term (ii) gives an upper bound on the probability of successful recovery from accessing exactly $r \in \{2, \dots, n - \ell\}$ incomplete nodes.

IV. PROBABILISTIC SYMMETRIC ALLOCATIONS

In the third variation of the storage allocation problem, we consider the case where each of the n nodes is selected by the source independently with probability $\min(\frac{\ell T}{n}, 1)$ to store $\frac{1}{\ell}$ amount of data, so that the expected total amount of storage used in the resulting *symmetric* allocation is at most $n \cdot \frac{\ell T}{n} \cdot \frac{1}{\ell} = T$, the given budget. The data collector subsequently accesses an r -subset of the n nodes selected uniformly at random from the collection of all $\binom{n}{r}$ possible r -subsets, where r is a given constant; successful recovery occurs if

and only if the total amount of data stored in the accessed nodes is at least 1. We seek an optimal probabilistic symmetric allocation of the budget T , specified by the value of parameter ℓ , that maximizes the probability of successful recovery, for a given choice of n , r , and T . Since successful recovery for a particular choice of ℓ occurs if and only if at least $\lceil 1 / (\frac{1}{\ell}) \rceil = \lceil \ell \rceil$ out of the r accessed nodes are nonempty, the corresponding probability of successful recovery can be written as

$$P_S(n, r, T, \ell) \triangleq \mathbb{P} \left[\mathcal{B} \left(r, \min \left(\frac{\ell T}{n}, 1 \right) \right) \geq \lceil \ell \rceil \right].$$

This optimization problem can therefore be expressed as follows:

$$\begin{aligned} \Pi_3(n, r, T) : \\ \underset{\ell}{\text{maximize}} \quad & \mathbb{P} \left[\mathcal{B} \left(r, \min \left(\frac{\ell T}{n}, 1 \right) \right) \geq \lceil \ell \rceil \right] \\ \text{subject to} \quad & \ell > 0. \end{aligned}$$

For budget $T \geq \frac{n}{r}$, the choice of $\ell = r$, which yields a recovery probability of $\mathbb{P}[\mathcal{B}(r, 1) \geq r] = 1$, is optimal.

Observe that the recovery probability $P_S(n, r, T, \ell)$ is zero when $\ell > r$. Furthermore, for fixed n , r , and T , the recovery probability is nondecreasing in ℓ within each of the unit intervals $(0, 1]$, $(1, 2]$, $(2, 3]$, \dots , since as ℓ increases within each interval, $\lceil \ell \rceil$ remains constant while $\min(\frac{\ell T}{n}, 1)$ either increases or remains constant at 1. Thus, given n , r , and T , we can find an optimal ℓ^* from among r candidates:

$$\{1, 2, \dots, r\}. \quad (21)$$

Fig. 7, which compares the performance of different probabilistic symmetric allocations over different budgets for an instance of r , suggests that there are two distinct phases pertaining to the optimal choice of ℓ : when the budget is below a certain threshold, the choice of $\ell = 1$, which corresponds to a minimal spreading of the budget (uncoded replication), is optimal; when the budget exceeds that same threshold, the choice of $\ell = r$, which corresponds to a maximal spreading of the budget, becomes optimal. This observation echoes our findings on the allocation and access models of the preceding sections, namely that minimal spreading ($\ell = 1$) is optimal for sufficiently small budgets, while maximal spreading ($\ell = r$) is optimal for sufficiently large budgets. However, we note two important distinctions in contrast to the previous models. First, the recovery probability for a probabilistic symmetric allocation in this model is a *continuous* nondecreasing function of the given budget; there are no ‘‘jumps’’ from one discrete value to the next. Second, our empirical computations suggest that

$$P_S^{\text{UB}} \triangleq \max_{\ell \in \{0, 1, \dots, \lfloor T \rfloor\}} \underbrace{1 - (1 - p)^\ell}_{\text{(i)}} + \underbrace{\mathbb{1}[\ell < \lfloor T \rfloor] \cdot (1 - p)^\ell}_{\text{(ii)}} + \sum_{r=2}^{n-\ell} \min \left(\underbrace{\frac{r(T-\ell)}{n-\ell}}_{\text{cf. Lemma 1}}, \underbrace{1 - \mathbb{1} \left[T - \ell < \frac{n-\ell}{r} \right] \cdot \frac{\gcd(r, r')}{\alpha \gcd(r, r') + r'}}_{\text{cf. Theorems 5 and 6}} \right) \cdot \mathbb{P}[\mathcal{B}(n-\ell, p) = r]. \quad (20)$$

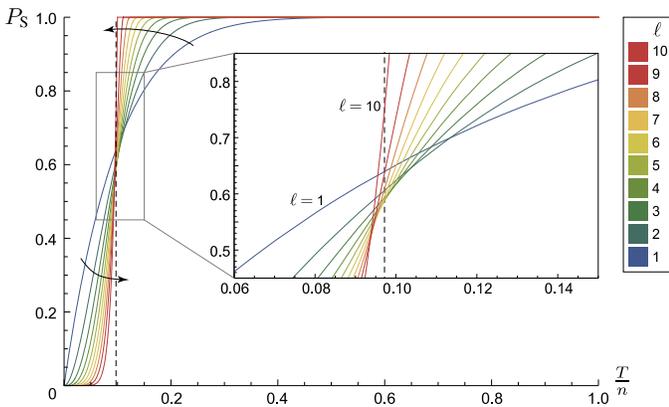


Fig. 7. Plot of recovery probability P_S against budget-per-node $\frac{T}{n}$ for each choice of parameter $\ell \in \{1, 2, \dots, r\}$, for $r = 10$. Parameter ℓ controls how much the budget is spread in the probabilistic symmetric allocation; specifically, each of the n nodes is selected by the source independently with probability $\min(\frac{\ell T}{n}, 1)$ to store $\frac{1}{\ell}$ amount of data. Arrows indicate the direction of increasing ℓ . The black dashed line marks the threshold on $\frac{T}{n}$ derived in Theorem 7; maximal spreading ($\ell = r$) is optimal for any $\frac{T}{n}$ greater than or equal to this threshold.

the phase transition from the optimality of minimal spreading to that of maximal spreading in this model is *sharp*; the other intermediate values of $\ell \in \{2, \dots, r-1\}$ never perform better than both $\ell = 1$ and $\ell = r$ simultaneously.

In Section IV-A, we shall demonstrate that the choice of $\ell = r$, which corresponds to a maximal spreading of the budget, is indeed optimal when the given budget T is sufficiently large, or equivalently, when a sufficiently high recovery probability is achievable.

A. Optimality of Maximal Spreading

Assume that $r \geq 2$. As noted earlier, the choice of $\ell = r$, which corresponds to a maximal spreading of the budget, is optimal for any $T \geq \frac{n}{r}$ because it yields the maximal recovery probability of 1. The following lemma provides an upper bound for the recovery probabilities corresponding to the *other* candidate values for ℓ^* in (21) at the critical budget $T = \frac{n}{r}$:

Lemma 6. *The probability of successful recovery $P_S(n, r, T, \ell)$ at $T = \frac{n}{r}$ is at most $\frac{3}{4}$ for any $\ell \in \{1, 2, \dots, r-1\}$.*

Such an upper bound allows us to derive a sufficient condition for the optimality of $\ell = r$, by making use of the fact that the recovery probability $P_S(n, r, T, \ell)$ is a nondecreasing function of the budget T . The following theorem shows that the choice of $\ell = r$ is optimal when the budget T is at least a specified threshold expressed in terms of n and r :

Theorem 7. *If*

$$T \geq \frac{n}{r} \left(\frac{3}{4} \right)^{\frac{1}{r}},$$

then the choice of $\ell = r$, which corresponds to a maximal spreading of the budget, is optimal.

The following corollary states an equivalent result in terms of the achievable recovery probability; it demonstrates the optimality of $\ell = r$ in the high recovery probability regime:

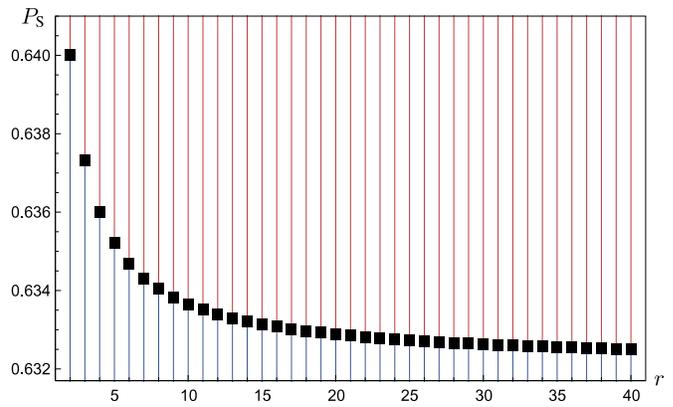


Fig. 8. Plot of recovery probability P_S against the number of nodes accessed r , indicating the value of P_S at which the optimal choice of parameter ℓ changes from 1 to r , for each given value of r . Specifically, if it is possible to achieve a recovery probability P_S above the square marker, then maximal spreading ($\ell = r$) is optimal; otherwise, minimal spreading or uncoded replication ($\ell = 1$) is optimal. Observe that the critical value of P_S for $r = 10$ (which is approximately 0.633652) corresponds to the intersection point of the curves for $\ell = 1$ and $\ell = 10$ in Fig. 7.

Corollary 4. *If a probability of successful recovery of at least $\frac{3}{4}$ is achievable for the given n , r , and T , then the choice of $\ell = r$ is optimal.*

Fig. 8 describes the optimal choice of ℓ for different values of r . We observe that the gap between the threshold of $\frac{3}{4}$ derived in Corollary 4 and the actual critical value of P_S indicated in the plot appears to be no more than 0.12.

V. CONCLUSION AND FUTURE WORK

We examined the problem of allocating a given total storage budget in a distributed storage system for maximum reliability. Three variations of the problem were studied in detail, and we are able to specify the optimal allocation or optimal symmetric allocation for a variety of cases. Although the exact optimal allocation is difficult to find in general, our results suggest a simple heuristic for achieving reliable storage: *when the budget is small, spread it minimally; when the budget is large, spread it maximally*. In other words, coding is unnecessary when the budget is small, but is beneficial when the budget is large.

The work in this paper can be extended in several directions. We can impose additional system design constraints on the model; one practical example is the application of a tighter per-node storage constraint $x_i \leq c_i < 1$. The independent probabilistic access model of Section II can be naturally generalized to the case of nonuniform access probabilities p_i for individual nodes. It would also be interesting to find reliable allocations for specific codes with desirable encoding or decoding properties, e.g., sparse codes that offer efficient algorithms (see, e.g., [24]–[27]). A related problem would be to construct such codes that work well under different allocations. Another set of interesting problems involves the application of richer access models; for instance, we can introduce a network topology to a set of storage nodes and data collectors, and allow each data collector to access only the nodes close to it. More generally, we can assign different priorities to each node for data storage and access, so as to

reflect the costs of storing data in the node and communicating with it.

APPENDIX PROOFS OF THEOREMS

Proof of Proposition 1: We note that the computational complexity of this problem was well understood in the Berkeley meetings [9] and is by no means a major contribution in this paper. We present the detailed proofs here for completeness.

Consider an allocation (x_1, \dots, x_n) where each x_i is a nonnegative rational number. The problem of computing the recovery probability of this allocation for the special case of $p = \frac{1}{2}$, for which $p^{|\mathbf{r}|}(1-p)^{n-|\mathbf{r}|} = (\frac{1}{2})^n$ for any subset $\mathbf{r} \subseteq \{1, \dots, n\}$, is equivalent to the counting version of the following decision problem (which happens to be polynomial-time solvable):

Definition. LARGEST SUBSET SUM (LSS)

Instance: Finite n -vector (a_1, \dots, a_n) with $a_i \in \mathbb{Z}_0^+$, and file size $d \in \mathbb{Z}^+$, where all a_i and d can be written as decimal numbers of length at most ℓ .

Question: Is there a subset $\mathbf{r} \subseteq \{1, \dots, n\}$ that satisfies $\sum_{i \in \mathbf{r}} a_i \geq d$?

Note that the allocation and file size have been scaled so that the problem parameters are all integers. We will proceed to show that the counting problem #LSS is #P-complete; this would in turn establish the #P-hardness of computing the recovery probability for an arbitrary value of p .

The index set \mathbf{r} can be represented as an n -vector of bits. Using this representation of \mathbf{r} as the certificate, it is easy to see that the binary relation corresponding to #LSS is both polynomially balanced (since the size of each certificate is n), and polynomial-time decidable (since the inequality can be verified in $O(n\ell)$ time for each certificate). It therefore follows that #LSS is in #P.

To show that #LSS is also #P-hard, we describe a polynomial-time Turing reduction of the #P-complete problem #3SAT [28] to #LSS. Our approach is similar to the standard method of reducing 3SAT to SUBSET SUM (see, e.g., [29]). Let ϕ be the Boolean formula in the given #3SAT instance; denote its m variables by v_1, \dots, v_m , and k clauses by C_1, \dots, C_k . To count the number of satisfying truth assignments for ϕ , we construct a #LSS instance with the help of Table III, whose entries are 0, 1, 2, or 3 (all blank entries are 0's). The entries of the n -vector for the #LSS instance are given by the first $(2m + 3k)$ rows of the table; the file size d is given by the last row of the table. Each entry a_i , $i \in \{1, \dots, 2m + 3k\}$, as well as d , is a positive integer with at most $(m + 2k)$ decimal digits. Observe that the set of satisfying truth assignments for ϕ can be put in a one-to-one correspondence with the collection of subsets $\mathbf{r} \subseteq \{1, \dots, 2m + 3k\}$ that satisfy $\sum_{i \in \mathbf{r}} a_i = d$; for each $i \in \{1, \dots, m\}$, we have " v_i " $\in \mathbf{r}$ if and only if $v_i = \text{TRUE}$, and " $\overline{v_i}$ " $\in \mathbf{r}$ if and only if $v_i = \text{FALSE}$. Therefore, if $f((a_1, \dots, a_n), d)$ is a subroutine for computing #LSS, then the number of satisfying truth assignments can

TABLE III
CONSTRUCTING A #LSS INSTANCE FOR A GIVEN #3SAT INSTANCE

	v_1	v_2	\dots	v_m	C_1		C_2		\dots	C_k	
v_1	1				$\mathbb{1}[v_1 \in C_1]$	0	$\mathbb{1}[v_1 \in C_2]$	0	\dots	$\mathbb{1}[v_1 \in C_k]$	0
$\overline{v_1}$	1				$\mathbb{1}[\overline{v_1} \in C_1]$	0	$\mathbb{1}[\overline{v_1} \in C_2]$	0	\dots	$\mathbb{1}[\overline{v_1} \in C_k]$	0
v_2		1			$\mathbb{1}[v_2 \in C_1]$	0	$\mathbb{1}[v_2 \in C_2]$	0	\dots	$\mathbb{1}[v_2 \in C_k]$	0
$\overline{v_2}$		1			$\mathbb{1}[\overline{v_2} \in C_1]$	0	$\mathbb{1}[\overline{v_2} \in C_2]$	0	\dots	$\mathbb{1}[\overline{v_2} \in C_k]$	0
\vdots					\vdots		\vdots		\vdots		
v_m				1	$\mathbb{1}[v_m \in C_1]$	0	$\mathbb{1}[v_m \in C_2]$	0	\dots	$\mathbb{1}[v_m \in C_k]$	0
$\overline{v_m}$				1	$\mathbb{1}[\overline{v_m} \in C_1]$	0	$\mathbb{1}[\overline{v_m} \in C_2]$	0	\dots	$\mathbb{1}[\overline{v_m} \in C_k]$	0
C_1					0	1					
					1	1					
					2	1					
C_2							0	1			
							1	1			
							2	1			
\vdots									\ddots		
C_k										0	1
										1	1
										2	1
d	1	1	\dots	1	3	1	3	1	\dots	3	1

be computed by calling f twice: first with d taking the value as prescribed above, and second with d taking the prescribed value *plus one*. The difference between the outputs from the two subroutine calls is equal to the number of distinct subsets \mathbf{r} that satisfy $\sum_{i \in \mathbf{r}} a_i = d$, which is equal to the number of satisfying truth assignments for ϕ . Finally, we note that this is indeed a polynomial-time Turing reduction since the table can be populated in $O(m^2k^2)$ simple steps, and the subroutine f is called exactly twice. ■

Proof of Lemma 1: Consider a feasible allocation (x_1, \dots, x_n) ; we have $\sum_{i=1}^n x_i \leq T$, where $x_i \geq 0$, $i = 1, \dots, n$. Let S_r denote the number of r -subsets of $\{x_1, \dots, x_n\}$ that have a sum of at least 1, where $r \in \{1, \dots, n\}$. By conditioning on the number of nodes accessed by the data collector, the probability of successful recovery for this allocation can be written as

$$\begin{aligned}
 & \mathbb{P}[\text{successful recovery}] \\
 &= \sum_{r=1}^n \mathbb{P}[\text{successful recovery} \mid \text{exactly } r \text{ nodes were accessed}] \cdot \\
 & \quad \mathbb{P}[\text{exactly } r \text{ nodes were accessed}] \\
 &= \sum_{r=1}^n \frac{S_r}{\binom{n}{r}} \cdot \mathbb{P}[\mathcal{B}(n, p) = r]. \tag{22}
 \end{aligned}$$

We proceed to find an upper bound for S_r . For a given r , we can write S_r inequalities of the form

$$x'_1 + \dots + x'_r \geq 1.$$

Summing up these S_r inequalities produces an inequality of the form

$$a_1x_1 + \dots + a_nx_n \geq S_r.$$

Since each x_i belongs to exactly $\binom{n-1}{r-1}$ distinct r -subsets of $\{x_1, \dots, x_n\}$, it follows that $0 \leq a_i \leq \binom{n-1}{r-1}$, $i = 1, \dots, n$.

Therefore,

$$S_r \leq a_1 x_1 + \cdots + a_n x_n \leq \binom{n-1}{r-1} \sum_{i=1}^n x_i \leq \binom{n-1}{r-1} T.$$

Since S_r is also at most $\binom{n}{r}$, i.e., the total number of r -subsets, we have

$$S_r \leq \min \left(\binom{n-1}{r-1} T, \binom{n}{r} \right).$$

Substituting this bound into (22) completes the proof. \blacksquare

Proof of Theorem 1: The suboptimality gap for the symmetric allocation $\bar{\mathbf{x}}(n, T, m=n)$ is at most the difference between its recovery probability (5) and the upper bound (6) from Lemma 1 for the optimal recovery probability. This difference is given by

$$\begin{aligned} & \sum_{r=1}^{\lceil \frac{n}{T} \rceil - 1} \frac{rT}{n} \binom{n}{r} p^r (1-p)^{n-r} \\ &= T \sum_{r=1}^{\lceil \frac{n}{T} \rceil - 1} \binom{n-1}{r-1} p^r (1-p)^{n-r} \\ &= pT \sum_{r=1}^{\lceil \frac{n}{T} \rceil - 1} \binom{n-1}{r-1} p^{r-1} (1-p)^{(n-1)-(r-1)} \\ &= pT \sum_{\ell=0}^{\lceil \frac{n}{T} \rceil - 2} \binom{n-1}{\ell} p^\ell (1-p)^{(n-1)-\ell} \\ &= pT \mathbb{P} \left[\mathcal{B}(n-1, p) \leq \left\lceil \frac{n}{T} \right\rceil - 2 \right] \triangleq \delta(n, p, T), \end{aligned}$$

as required. Assuming now that $p > \frac{1}{T}$, we have

$$\begin{aligned} \delta(n, p, T) &\leq pT \mathbb{P} \left[\mathcal{B}(n-1, p) \leq \frac{n-1}{T} \right] \\ &= pT \mathbb{P} \left[\mathcal{B}(n-1, p) \leq \frac{1}{pT} (n-1)p \right] \\ &\leq pT \exp \left(-\frac{(n-1)p}{2} \left(1 - \frac{1}{pT} \right)^2 \right). \end{aligned} \quad (23)$$

Inequality (23) follows from the fact that

$$\left\lceil \frac{n}{T} \right\rceil - 2 < \frac{n}{T} + 1 - 2 < \frac{n}{T} - \frac{1}{T}.$$

Inequality (24) follows from the observation that $\frac{1}{pT} \in (0, 1)$, and the subsequent application of the Chernoff bound for deviation below the mean of the binomial distribution (see, e.g., [30]). For fixed p and T , this upper bound approaches zero as n goes to infinity. \blacksquare

Proof of Theorem 2: We compare the recovery probability of $\bar{\mathbf{x}}(n, T, m=\lceil T \rceil)$ to an upper bound for the recovery probabilities of all other allocations.

Suppose that $1 < T < n$. Recall from (7) that the probability of successful recovery for $\bar{\mathbf{x}}(n, T, m=\lceil T \rceil)$ is given by

$$P_1(p, T) \triangleq 1 - (1-p)^{\lceil T \rceil}.$$

Consider a feasible allocation (x_1, \dots, x_n) ; we have $\sum_{i=1}^n x_i \leq T$, where $x_i \geq 0$, $i = 1, \dots, n$. Let ℓ be the number of individual nodes in this allocation that store at least

a unit amount of data; for brevity, we refer to these nodes as being *complete*. It follows from the budget constraint that the number of complete nodes $\ell \in \{0, 1, \dots, \lceil T \rceil\}$. When $\ell = \lceil T \rceil$, the allocation has a recovery probability identical to $P_1(p, T)$. Now, assuming that $\ell \in \{0, 1, \dots, \lceil T \rceil - 1\}$, successful recovery can occur in two ways:

- (i) when the accessed subset contains one or more complete nodes, which occurs with probability $1 - (1-p)^\ell$, or
- (ii) when the accessed subset contains no complete nodes but has a sum of at least 1.

In case (ii), the accessed subset would consist of two or more *incomplete* nodes. Using the argument in the proof of Lemma 1, we can show that there are at most

$$\min \left(\binom{n-\ell-1}{r-1} (T-\ell), \binom{n-\ell}{r} \right)$$

r -subsets of incomplete nodes whose sum is at least 1, since the total amount of data stored over the $n-\ell$ incomplete nodes is at most $T-\ell$. It follows then that the recovery probability for a feasible allocation with exactly $\ell \in \{0, 1, \dots, \lceil T \rceil - 1\}$ complete nodes is at most

$$P_2(n, p, T, \ell) \triangleq 1 - (1-p)^\ell + (1-p)^\ell \sum_{r=2}^{n-\ell} \min \left(\frac{T-\ell}{n-\ell} \cdot r, 1 \right) \binom{n-\ell}{r} p^r (1-p)^{n-\ell-r}.$$

Thus,

$$P_1(p, T) \geq P_2(n, p, T, \ell)$$

for all $\ell \in \{0, 1, \dots, \lceil T \rceil - 1\}$ is a sufficient condition for $\bar{\mathbf{x}}(n, T, m=\lceil T \rceil)$ to be an optimal allocation. After further simplification of this inequality, we arrive at inequality (8) as required. \blacksquare

Proof of Corollary 1: Suppose that $1 < T < n$. We will show that the sufficient condition of Theorem 2 is satisfied for any $p \leq \frac{2}{(n-\lceil T \rceil)^2}$. Note that when $n - \lceil T \rceil = 1$, or equivalently $T \in [n-1, n)$, we have to show that $\bar{\mathbf{x}}(n, T, m=\lceil T \rceil)$ is an optimal allocation for *any* p , i.e., in the interval $(0, 1)$.

First, observe that the summation term in inequality (8) is always nonnegative, i.e.,

$$\sum_{r=2}^{\lceil \frac{n-\ell}{T-\ell} \rceil - 1} \left(1 - \frac{T-\ell}{n-\ell} \cdot r \right) \binom{n-\ell}{r} \left(\frac{p}{1-p} \right)^r \geq 0,$$

since for any $r \in \{2, \dots, \lceil \frac{n-\ell}{T-\ell} \rceil - 1\}$ and $\ell \in \{0, 1, \dots, \lceil T \rceil - 1\}$, we have

$$r \leq \left\lceil \frac{n-\ell}{T-\ell} \right\rceil - 1 \iff r < \frac{n-\ell}{T-\ell} \iff 1 - \frac{T-\ell}{n-\ell} \cdot r > 0.$$

Therefore, a simpler but weaker sufficient condition for $\bar{\mathbf{x}}(n, T, m=\lceil T \rceil)$ to be an optimal allocation is

$$\begin{aligned} & 1 - (1-p)^{\lceil T \rceil - n} + (n - (\lceil T \rceil - 1)) \left(\frac{p}{1-p} \right) \geq 0 \\ & \iff 1 + (n - \lceil T \rceil)p - (1-p)^{1-(n-\lceil T \rceil)} \geq 0, \end{aligned}$$

which is an inequality in only two variables p and $s \triangleq n - \lceil T \rceil$, where $s \in \{1, \dots, n-1\}$. When $s = 1$, or

equivalently $T \in [n-1, n)$, this inequality is satisfied for any $p \in (0, 1)$, as required. Defining the function

$$f(s, p) \triangleq 1 + sp - (1-p)^{1-s},$$

it suffices to show that $f(s, p) \geq 0$ for any $s \in \mathbb{Z}^+$, $s \geq 2$, and $p \in (0, \frac{2}{s^2}]$. We do this by demonstrating that for any $s \in \mathbb{Z}^+$, $s \geq 2$, the function $f(s, p)$ is concave in p on the interval $p \in (0, \frac{2}{s^2}]$, and is nonnegative at both endpoints, i.e., $f(s, p=0) \geq 0$ and $f(s, p=\frac{2}{s^2}) \geq 0$.

The second-order partial derivative of $f(s, p)$ wrt p is given by

$$\frac{\partial^2}{\partial p^2} f(s, p) = -s(s-1)(1-p)^{-1-s}.$$

Since $\frac{\partial^2}{\partial p^2} f(s, p) < 0$ for any $s \in \mathbb{Z}^+$, $s \geq 2$, and $p \in (0, \frac{2}{s^2}]$, it follows that the function $f(s, p)$ is concave in p on the interval $p \in (0, \frac{2}{s^2}]$ for any $s \in \mathbb{Z}^+$, $s \geq 2$.

Suppose that $s \in \mathbb{Z}^+$, $s \geq 2$. Clearly, $f(s, p=0) = 0$. To show that $f(s, p=\frac{2}{s^2}) \geq 0$, we define the function

$$g(s) \triangleq \ln\left(1 + \frac{2}{s}\right) + (s-1) \ln\left(1 - \frac{2}{s^2}\right),$$

and show that $g(s) \geq 0$ for any $s \in \mathbb{Z}^+$, $s \geq 2$. Direct evaluation of the function gives us $g(s=2) = 0$, and $g(s=3) = \ln \frac{5}{3} - 2 \ln \frac{9}{7} > 0$. For $s \geq 4$, we consider the derivatives of $g(s)$:

$$g'(s) = \frac{1}{s} + \frac{1}{s+2} - \frac{2(s-2)}{s^2-2} + \ln\left(1 - \frac{2}{s^2}\right),$$

$$g''(s) = \frac{8(s^3 - s^2 - 6s - 2)}{s^2(s+2)^2(s^2-2)^2}.$$

Since $g''(s) \geq 0$ for any $s \geq 4$, and $\lim_{s \rightarrow \infty} g'(s) = 0$, it follows that $g'(s) \leq 0$ for any $s \geq 4$. Now, since $g'(s) \leq 0$ for any $s \geq 4$, and $\lim_{s \rightarrow \infty} g(s) = 0$, it follows that $g(s) \geq 0$ for any $s \geq 4$. Therefore, for any $s \in \mathbb{Z}^+$, $s \geq 2$, we have

$$\ln\left(1 + \frac{2}{s}\right) + (s-1) \ln\left(1 - \frac{2}{s^2}\right) = g(s) \geq 0$$

$$\iff 1 + \frac{2}{s} \geq \left(1 - \frac{2}{s^2}\right)^{1-s}$$

$$\iff f\left(s, p=\frac{2}{s^2}\right) \geq 0,$$

as required. \blacksquare

Proof of Theorem 3: We will show that if condition (11) is satisfied, then $\Delta(p, T, k) \geq 0$ for any $k \in \mathbb{Z}^+$. First, we note that

$$\frac{\binom{\lfloor kT \rfloor}{k-1}}{\binom{\lfloor kT \rfloor}{k}} = \frac{k}{\lfloor kT \rfloor - k + 1}$$

$$= \frac{k}{\lfloor k(\lfloor T \rfloor + \tau) \rfloor - k + 1}, \text{ where } \tau \triangleq T - \lfloor T \rfloor \in [0, 1)$$

$$= \frac{k}{k\lfloor T \rfloor + \lfloor k\tau \rfloor - k + 1}$$

$$\geq \frac{k}{k\lfloor T \rfloor} \quad (25)$$

$$= \frac{1}{\lfloor T \rfloor}. \quad (26)$$

Inequality (25) follows from the fact that

$$\lfloor k\tau \rfloor \leq k\tau < k \iff \lfloor k\tau \rfloor \leq k-1 \iff \lfloor k\tau \rfloor - k + 1 \leq 0.$$

Now, if condition (11) is satisfied, then we necessarily have $T \geq 2$; otherwise, $T \in [1, 2)$ would imply that $\lfloor T \rfloor = 1$, which produces $(1-p)^{\lfloor T \rfloor} + 2\lfloor T \rfloor p(1-p)^{\lfloor T \rfloor - 1} - 1 = p > 0$, contradicting our assumption. It follows that

$$(1-p)^{\lfloor T \rfloor} + 2\lfloor T \rfloor p(1-p)^{\lfloor T \rfloor - 1} - 1 \leq 0$$

$$\iff \mathbb{P}[\mathcal{B}(\lfloor T \rfloor, p) = 0] + 2\mathbb{P}[\mathcal{B}(\lfloor T \rfloor, p) = 1] - 1 \leq 0$$

$$\iff \mathbb{P}[\mathcal{B}(\lfloor T \rfloor, p) \geq 2] \geq \mathbb{P}[\mathcal{B}(\lfloor T \rfloor, p) = 1]$$

$$\iff \sum_{j=2}^{\lfloor T \rfloor} \binom{\lfloor T \rfloor}{j} p^j (1-p)^{\lfloor T \rfloor - j} \geq \lfloor T \rfloor p(1-p)^{\lfloor T \rfloor - 1}$$

$$\iff \sum_{j=2}^{\lfloor T \rfloor} \frac{1}{\lfloor T \rfloor} \binom{\lfloor T \rfloor}{j} \left(\frac{p}{1-p}\right)^{j-1} \geq 1 \quad (27)$$

$$\implies \sum_{j=2}^{\lfloor T \rfloor} \frac{1}{\lfloor T \rfloor} \binom{\lfloor T \rfloor}{j} \left(\frac{p}{1-p}\right)^{j-1} \geq 1. \quad (28)$$

Observe that $\alpha_{k,T} \triangleq \lfloor (k+1)T \rfloor - \lfloor kT \rfloor \in \{\lfloor T \rfloor, \lceil T \rceil\}$, because $\alpha_{k,T} \in (T-1, T+1)$ and there are only two integers $\lfloor T \rfloor$ and $\lceil T \rceil$, which are possibly nondistinct, in this interval. It follows from (27) and (28) that

$$\sum_{j=2}^{\alpha_{k,T}} \frac{1}{\lfloor T \rfloor} \binom{\alpha_{k,T}}{j} \left(\frac{p}{1-p}\right)^{j-1} \geq 1. \quad (29)$$

Therefore, we have

$$\min(\alpha_{k,T}-1, k) \sum_{i=1}^{\alpha_{k,T}} \sum_{j=i+1}^{\lfloor kT \rfloor} \frac{\binom{\lfloor kT \rfloor}{k-i}}{\binom{\lfloor kT \rfloor}{k}} \left(\alpha_{k,T}\right) \left(\frac{p}{1-p}\right)^{-i+j}$$

$$\geq \sum_{i=1}^1 \sum_{j=i+1}^{\lfloor kT \rfloor} \frac{\binom{\lfloor kT \rfloor}{k-i}}{\binom{\lfloor kT \rfloor}{k}} \left(\alpha_{k,T}\right) \left(\frac{p}{1-p}\right)^{-i+j} \quad (30)$$

$$= \sum_{j=2}^{\alpha_{k,T}} \frac{\binom{\lfloor kT \rfloor}{k-1}}{\binom{\lfloor kT \rfloor}{k}} \left(\alpha_{k,T}\right) \left(\frac{p}{1-p}\right)^{j-1}$$

$$\geq \sum_{j=2}^{\alpha_{k,T}} \frac{1}{\lfloor T \rfloor} \binom{\alpha_{k,T}}{j} \left(\frac{p}{1-p}\right)^{j-1}, \text{ from (26)}$$

$$\geq 1, \text{ from (29).}$$

Inequality (30) follows from the fact that

$$\min(\alpha_{k,T}-1, k) \geq \min(2-1, 1) = 1.$$

Consequently,

$$\min(\alpha_{k,T}-1, k) \sum_{i=1}^{\alpha_{k,T}} \sum_{j=i+1}^{\lfloor kT \rfloor} \frac{\binom{\lfloor kT \rfloor}{k-i}}{\binom{\lfloor kT \rfloor}{k}} \left(\alpha_{k,T}\right) \left(\frac{p}{1-p}\right)^{-i+j} \geq \binom{\lfloor kT \rfloor}{k}$$

$$\iff \Delta(p, T, k) \geq 0, \text{ from (10).}$$

It follows that

$$\begin{aligned} P_S(p, T, m=\lfloor T \rfloor) &\leq P_S(p, T, m=\lfloor 2T \rfloor) \\ &\leq \dots \leq P_S(p, T, m=\lfloor \lfloor \frac{n}{T} \rfloor T \rfloor), \end{aligned}$$

and so we conclude that an optimal m^* is given by either $m = \lfloor \lfloor \frac{n}{T} \rfloor T \rfloor$ or $m = n$. \blacksquare

Proof of Corollary 2: If $p \geq \frac{4}{3\lfloor T \rfloor}$, then we necessarily have $T \geq 2$; otherwise, $T \in [1, 2)$ would imply that $\lfloor T \rfloor = 1$, which produces $p \geq \frac{4}{3\lfloor T \rfloor} = \frac{4}{3}$, contradicting the definition of p . We will show that condition (11) of Theorem 3 is satisfied for any $T \geq 2$ and $p \geq \frac{4}{3\lfloor T \rfloor}$. To do this, we define the function

$$f(p, T) \triangleq (1-p)^{\lfloor T \rfloor} + 2\lfloor T \rfloor p(1-p)^{\lfloor T \rfloor - 1} - 1,$$

and show that $f(p, T) \leq f(p = \frac{4}{3\lfloor T \rfloor}, T) \leq 0$ for any $T \geq 2$ and $p \geq \frac{4}{3\lfloor T \rfloor}$.

The partial derivative of $f(p, T)$ wrt p is given by

$$\frac{\partial}{\partial p} f(p, T) = \lfloor T \rfloor (1-p)^{\lfloor T \rfloor - 2} (1+p-2\lfloor T \rfloor p).$$

Observe that $f(p, T)$ is decreasing wrt p for any $T \geq 2$ and $p \geq \frac{4}{3\lfloor T \rfloor}$, since

$$p \geq \frac{4}{3\lfloor T \rfloor} = \frac{1}{\frac{3}{4}\lfloor T \rfloor} > \frac{1}{2\lfloor T \rfloor - 1}$$

$$\implies 2\lfloor T \rfloor p - p > 1 \iff 1 + p - 2\lfloor T \rfloor p < 0 \iff \frac{\partial}{\partial p} f(p, T) < 0.$$

Now, consider the function

$$g(T) \triangleq f(p = \frac{4}{3\lfloor T \rfloor}, T) = \left(1 - \frac{4}{3\lfloor T \rfloor}\right)^{\lfloor T \rfloor - 1} \left(\frac{11}{3} - \frac{4}{3\lfloor T \rfloor}\right) - 1.$$

We will proceed to show that $g(T) \leq 0$ for any $T \geq 2$. For $T \in [2, 3)$, we have $\lfloor T \rfloor = 2$ and $g(T) = 0$. To show that $g(T) \leq 0$ for any $T \geq 3$, we consider the function

$$h(T) \triangleq (T-1) \ln \left(1 - \frac{4}{3T}\right) + \ln \left(\frac{11}{3} - \frac{4}{3T}\right),$$

which has the derivatives

$$h'(T) = \frac{1}{3T-4} + \frac{11}{11T-4} + \ln \left(1 - \frac{4}{3T}\right),$$

$$h''(T) = \frac{16(11T^2 - 24T - 16)}{T(33T^2 - 56T + 16)^2}.$$

Since $h''(T) > 0$ for any $T \geq 3$, and $\lim_{T \rightarrow \infty} h'(T) = 0$, it follows that $h'(T) \leq 0$ for any $T \geq 3$. Now, since $h'(T) \leq 0$ for any $T \geq 3$, and $h(T=3) = \ln \frac{29}{9} - 2 \ln \frac{9}{5} < 0$, it follows that $h(T) < 0$ for any $T \geq 3$. Thus, for any $T \geq 3$, we have

$$(\lfloor T \rfloor - 1) \ln \left(1 - \frac{4}{3\lfloor T \rfloor}\right) + \ln \left(\frac{11}{3} - \frac{4}{3\lfloor T \rfloor}\right) = h(\lfloor T \rfloor) < 0$$

$$\iff \ln \left\{ \left(1 - \frac{4}{3\lfloor T \rfloor}\right)^{\lfloor T \rfloor - 1} \left(\frac{11}{3} - \frac{4}{3\lfloor T \rfloor}\right) \right\} < 0$$

$$\iff \left(1 - \frac{4}{3\lfloor T \rfloor}\right)^{\lfloor T \rfloor - 1} \left(\frac{11}{3} - \frac{4}{3\lfloor T \rfloor}\right) < 1 \iff g(T) < 0.$$

Combining these results, we obtain

$$f(p, T) \leq f\left(p = \frac{4}{3\lfloor T \rfloor}, T\right) = g(T) \leq 0$$

for any $T \geq 2$ and $p \geq \frac{4}{3\lfloor T \rfloor}$, as required. \blacksquare

Proof of Lemma 2: Suppose that $T > 1$. We will show that if condition (12) or condition (13) is satisfied, then $\Delta(p, T, k) \leq 0$ for any $k \in \mathbb{Z}^+$. First, we note that for any $i \in \{1, \dots, k\}$,

$$\begin{aligned} \frac{\binom{\lfloor kT \rfloor}{k-i}}{\binom{\lfloor kT \rfloor}{k}} &= \frac{\overbrace{(k)(k-1)\dots(k-i+1)}^{i \text{ terms}}}{\underbrace{(\lfloor kT \rfloor - k + i)\dots(\lfloor kT \rfloor - k + 2)(\lfloor kT \rfloor - k + 1)}_{i \text{ terms}}} \\ &\leq \left(\frac{k}{\lfloor kT \rfloor - k + 1}\right)^i \\ &\leq \left(\frac{k}{kT - 1 - k + 1}\right)^i \\ &= \left(\frac{1}{T-1}\right)^i. \end{aligned} \quad (31)$$

Now, if condition (12) is satisfied, then

$$\begin{aligned} &\sum_{i=1}^{\lfloor T \rfloor - 1} \sum_{j=i+1}^{\lfloor T \rfloor} \left(\frac{1}{T-1}\right)^i \binom{\lfloor T \rfloor}{j} \left(\frac{p}{1-p}\right)^{-i+j} \\ &= \sum_{i=1}^{T-1} \sum_{j=i+1}^T \left(\frac{1}{T-1}\right)^i \binom{T}{j} \left(\frac{1}{1-\frac{1}{T}}\right)^{-i+j} \\ &= \sum_{i=1}^{T-1} \sum_{j=i+1}^T \binom{T}{j} \left(\frac{1}{T-1}\right)^j \\ &= \sum_{\ell=2}^T (\ell-1) \binom{T}{\ell} \left(\frac{1}{T-1}\right)^\ell = 1. \end{aligned}$$

On the other hand, if condition (13) is satisfied, then

$$\begin{aligned} &\sum_{i=1}^{\lfloor T \rfloor - 1} \sum_{j=i+1}^{\lfloor T \rfloor} \left(\frac{1}{T-1}\right)^i \binom{\lfloor T \rfloor}{j} \left(\frac{p}{1-p}\right)^{-i+j} \\ &= \sum_{i=1}^{\lfloor T \rfloor - 1} \sum_{j=i+1}^{\lfloor T \rfloor} \binom{\lfloor T \rfloor}{j} \left(\frac{1-p}{p(T-1)}\right)^i \left(\frac{p}{1-p}\right)^j \\ &= \sum_{\ell=2}^{\lfloor T \rfloor} \left(\sum_{r=1}^{\ell-1} \left(\frac{1-p}{p(T-1)}\right)^r\right) \binom{\lfloor T \rfloor}{\ell} \left(\frac{p}{1-p}\right)^\ell \\ &= 1 - \frac{T \left(\frac{1}{T} (1 - \frac{1}{T})^{\lfloor T \rfloor - 1} - p(1-p)^{\lfloor T \rfloor - 1}\right)}{(1-pT)(1 - \frac{1}{T})^{\lfloor T \rfloor - 1} (1-p)^{\lfloor T \rfloor - 1}} \leq 1. \end{aligned}$$

Thus, if either condition is satisfied, we have

$$\begin{aligned} &\sum_{i=1}^{\lfloor T \rfloor - 1} \sum_{j=i+1}^{\lfloor T \rfloor} \left(\frac{1}{T-1}\right)^i \binom{\lfloor T \rfloor}{j} \left(\frac{p}{1-p}\right)^{-i+j} \leq 1 \quad (32) \\ \implies &\sum_{i=1}^{\lfloor T \rfloor - 1} \sum_{j=i+1}^{\lfloor T \rfloor} \left(\frac{1}{T-1}\right)^i \binom{\lfloor T \rfloor}{j} \left(\frac{p}{1-p}\right)^{-i+j} \leq 1. \quad (33) \end{aligned}$$

As in the proof of Theorem 3, we note that $\alpha_{k,T} \triangleq \lfloor (k+1)T \rfloor - \lfloor kT \rfloor \in \{\lfloor T \rfloor, \lfloor T \rfloor\}$. It follows from (32) and (33) that

$$\sum_{i=1}^{\alpha_{k,T}-1} \sum_{j=i+1}^{\alpha_{k,T}} \left(\frac{1}{T-1}\right)^i \binom{\alpha_{k,T}}{j} \left(\frac{p}{1-p}\right)^{-i+j} \leq 1. \quad (34)$$

Therefore, we have

$$\begin{aligned}
& \min_{i=1}^{\alpha_{k,T}-1,k} \sum_{j=i+1}^{\alpha_{k,T}} \frac{\binom{\lfloor kT \rfloor}{k-i}}{\binom{\lfloor kT \rfloor}{k}} \left(\alpha_{k,T} \right) \left(\frac{p}{1-p} \right)^{-i+j} \\
& \leq \sum_{i=1}^{\min(\alpha_{k,T}-1,k)} \sum_{j=i+1}^{\alpha_{k,T}} \left(\frac{1}{T-1} \right)^i \left(\alpha_{k,T} \right) \left(\frac{p}{1-p} \right)^{-i+j}, \text{ from (31)} \\
& \leq \sum_{i=1}^{\alpha_{k,T}-1} \sum_{j=i+1}^{\alpha_{k,T}} \left(\frac{1}{T-1} \right)^i \left(\alpha_{k,T} \right) \left(\frac{p}{1-p} \right)^{-i+j} \\
& \leq 1, \text{ from (34)}.
\end{aligned}$$

Consequently,

$$\begin{aligned}
& \min_{i=1}^{\alpha_{k,T}-1,k} \sum_{j=i+1}^{\alpha_{k,T}} \binom{\lfloor kT \rfloor}{k-i} \left(\alpha_{k,T} \right) \left(\frac{p}{1-p} \right)^{-i+j} \leq \binom{\lfloor kT \rfloor}{k} \\
& \iff \Delta(p, T, k) \leq 0, \text{ from (10)}.
\end{aligned}$$

It follows that

$$\begin{aligned}
P_S(p, T, m=\lfloor T \rfloor) & \geq P_S(p, T, m=\lfloor 2T \rfloor) \\
& \geq P_S(p, T, m=\lfloor 3T \rfloor) \geq \dots,
\end{aligned}$$

and since

$$P_S(p, T, m=n) \begin{cases} = P_S(p, T, m=\lfloor \frac{n}{T} \rfloor T) & \text{if } \frac{n}{T} \in \mathbb{Z}^+, \\ \leq P_S(p, T, m=\lfloor (\lfloor \frac{n}{T} \rfloor + 1) T \rfloor) & \text{otherwise,} \end{cases}$$

we conclude that an optimal m^* is given by $m = \lfloor T \rfloor$. \blacksquare

Proof of Lemma 3: Since $\bar{x}(n, T, m=\lfloor T \rfloor)$ is indeed optimal for any p when $T = 1$, we need only consider the case of $T > 1$. We will show that either condition (12) or condition (13) of Lemma 2 is satisfied for any $T > 1$ and $p \leq \frac{2}{\lfloor T \rfloor} - \frac{1}{T}$. We do this in two steps: First, we define the function

$$f(p, T) \triangleq \frac{p(1-p)^{\lfloor T \rfloor - 1}}{\frac{1}{T} \left(1 - \frac{1}{T} \right)^{\lfloor T \rfloor - 1}} - 1,$$

and show that $f(p, T) \leq f\left(p = \frac{2}{\lfloor T \rfloor} - \frac{1}{T}, T\right) \leq 0$ for any $T > 1$ and $p \leq \frac{2}{\lfloor T \rfloor} - \frac{1}{T}$. Second, we apply the appropriate condition from Lemma 2 for each pair of T and p .

The partial derivative of $f(p, T)$ wrt p is given by

$$\frac{\partial}{\partial p} f(p, T) = \frac{(1-p)^{\lfloor T \rfloor} (1-p)^{\lfloor T \rfloor - 2}}{\frac{1}{T} \left(1 - \frac{1}{T} \right)^{\lfloor T \rfloor - 1}}.$$

Observe that $f(p, T)$ is nondecreasing wrt p for any $T > 1$ and $p \leq \frac{2}{\lfloor T \rfloor} - \frac{1}{T}$, since

$$\begin{aligned}
p \leq \frac{2}{\lfloor T \rfloor} - \frac{1}{T} & \leq \frac{2}{\lfloor T \rfloor} - \frac{1}{\lfloor T \rfloor} = \frac{1}{\lfloor T \rfloor} \\
\implies p \lfloor T \rfloor \leq 1 & \iff 1 - p \lfloor T \rfloor \geq 0 \iff \frac{\partial}{\partial p} f(p, T) \geq 0.
\end{aligned}$$

Now, consider the function

$$g(T) \triangleq f\left(p = \frac{2}{\lfloor T \rfloor} - \frac{1}{T}, T\right) = \frac{\left(\frac{2}{\lfloor T \rfloor} - \frac{1}{T}\right) \left(1 - \frac{2}{\lfloor T \rfloor} + \frac{1}{T}\right)^{\lfloor T \rfloor - 1}}{\frac{1}{T} \left(1 - \frac{1}{T}\right)^{\lfloor T \rfloor - 1}} - 1.$$

We will proceed to show that $g(T) \leq 0$ for any $T > 1$ by reparameterizing $g(T)$ as $h(c, \tau)$, where $c \triangleq \lfloor T \rfloor$ and $\tau \triangleq \lfloor T \rfloor - T$:

$$h(c, \tau) \triangleq g(T=c-\tau) = \frac{\left(\frac{2}{c} - \frac{1}{c-\tau}\right) \left(1 - \frac{2}{c} + \frac{1}{c-\tau}\right)^{c-1}}{\frac{1}{c-\tau} \left(1 - \frac{1}{c-\tau}\right)^{c-1}} - 1.$$

The partial derivative of $h(c, \tau)$ wrt τ is given by

$$\frac{\partial}{\partial \tau} h(c, \tau) = -\frac{2\tau^2(c-2) \left(1 - \frac{2}{c} + \frac{1}{c-\tau}\right)^c}{(c(c-1-\tau) + 2\tau)^2 \left(1 - \frac{1}{c-\tau}\right)^c}.$$

Since $\frac{\partial}{\partial \tau} h(c, \tau) \leq 0$ for any $c \in \mathbb{Z}^+$, $c \geq 2$, and $\tau \in [0, 1)$, it follows that for any $T > 1$, we have

$$\begin{aligned}
g(T) & = h(c=\lfloor T \rfloor, \tau=\lfloor T \rfloor - T) \\
& \leq h(c=\lfloor T \rfloor, \tau=0) \\
& = \frac{\left(\frac{2}{\lfloor T \rfloor} - \frac{1}{\lfloor T \rfloor}\right) \left(1 - \frac{2}{\lfloor T \rfloor} + \frac{1}{\lfloor T \rfloor}\right)^{\lfloor T \rfloor - 1}}{\frac{1}{\lfloor T \rfloor} \left(1 - \frac{1}{\lfloor T \rfloor}\right)^{\lfloor T \rfloor - 1}} - 1 = 0.
\end{aligned}$$

Combining these results, we obtain

$$f(p, T) \leq f\left(p = \frac{2}{\lfloor T \rfloor} - \frac{1}{T}, T\right) = g(T) \leq 0$$

for any $T > 1$ and $p \leq \frac{2}{\lfloor T \rfloor} - \frac{1}{T}$, which implies

$$p(1-p)^{\lfloor T \rfloor - 1} \leq \frac{1}{T} \left(1 - \frac{1}{T}\right)^{\lfloor T \rfloor - 1}.$$

Finally, we apply the appropriate condition from Lemma 2 for each pair of T and p . For $T \in \mathbb{Z}^+$, $T > 1$, we have $\frac{2}{\lfloor T \rfloor} - \frac{1}{T} = \frac{1}{T}$: we use condition (12) for $p = \frac{1}{T}$, and condition (13) for $p < \frac{1}{T}$. For $T \notin \mathbb{Z}^+$, $T > 1$, we have $\frac{2}{\lfloor T \rfloor} - \frac{1}{T} < \frac{1}{T}$: we use condition (13) for $p < \frac{1}{T}$. \blacksquare

Proof of Theorem 4: Since $\bar{x}(n, T, m=\lfloor T \rfloor)$ is indeed optimal for any p when $T = 1$, we need only consider the case of $T > 1$. We will show that $\bar{x}(n, T, m=\lfloor T \rfloor)$ is an optimal symmetric allocation for any $T > 1$ and $p \leq \frac{1}{\lfloor T \rfloor}$. We do this by considering subintervals of T over which $\lfloor T \rfloor$ is constant.

Let T be confined to the unit interval $(c, c+1]$, where $c \in \mathbb{Z}^+$. According to Lemma 3, $\bar{x}(n, T, m=\lfloor T \rfloor)$ is optimal for any $p \in \left(0, \frac{2}{c+1} - \frac{1}{T}\right]$ and $T \in (c, c+1]$, or equivalently, for any

$$p \in \left(0, \frac{1}{c+1}\right] \text{ and } T \in \left[\frac{1}{\frac{2}{c+1} - p}, c+1\right] \cap (c, c+1].$$

This is just the area below a ‘‘peak’’ in Fig. 4, expressed in terms of different independent variables. For each $p \in \left(0, \frac{1}{c+1}\right)$, we can always find a T_0 such that

$$T_0 \in \left[\frac{1}{\frac{2}{c+1} - p}, c+1\right] \cap (c, c+1).$$

For example, we can pick $T_0 = c+1 - \delta$, where

$$\delta \triangleq \frac{1}{2} \left(c+1 - \max\left(c, \frac{1}{\frac{2}{c+1} - p}\right) \right) \in (0, 1).$$

Now, we make the crucial observation that if $\bar{x}(n, T, m=\lfloor T \rfloor)$ is an optimal symmetric allocation for $T = T_0$, then $\bar{x}(n, T, m=\lfloor T \rfloor)$ is also an optimal symmetric allocation for any $T \in [\lfloor T_0 \rfloor, T_0]$. This claim can be proven by contradiction: the recovery probability for $\bar{x}(n, T, m=\lfloor T \rfloor)$ is given by

$$P_S(p, T, m=\lfloor T \rfloor) = \mathbb{P}[\mathcal{B}(\lfloor T \rfloor, p) \geq 1]$$

which remains constant for all $T \in [\lfloor T_0 \rfloor, T_0]$, and a symmetric allocation that performs strictly better than $\bar{x}(n, T, m=\lfloor T \rfloor)$ for some $T \in [\lfloor T_0 \rfloor, T_0]$ would therefore also outperform $\bar{x}(n, T, m=\lfloor T \rfloor)$ for $T = T_0$. Since $\bar{x}(n, T, m=\lfloor T \rfloor)$ is indeed optimal for our choice of T_0 , it follows then that $\bar{x}(n, T, m=\lfloor T \rfloor)$ is also optimal for any

$$p \in \left(0, \frac{1}{c+1}\right) \text{ and } T \in (c, c+1].$$

By applying this result for each $c \in \mathbb{Z}^+$, we reach the conclusion that $\bar{x}(n, T, m=\lfloor T \rfloor)$ is an optimal symmetric allocation for any $T > 1$ and $p < \frac{1}{\lfloor T \rfloor}$.

Finally, to extend the optimality of $\bar{x}(n, T, m=\lfloor T \rfloor)$ to $p = \frac{1}{\lfloor T \rfloor}$, we note that the recovery probability $P_S(p, T, m) \triangleq \mathbb{P}[\mathcal{B}(m, p) \geq \frac{m}{T}]$ is a polynomial in p and is therefore continuous at $p = \frac{1}{\lfloor T \rfloor}$. Since $\bar{x}(n, T, m=\lfloor T \rfloor)$ is optimal as $p \rightarrow \frac{1}{\lfloor T \rfloor}^-$, it remains optimal at $p = \frac{1}{\lfloor T \rfloor}$. ■

Proof of Proposition 2: Consider an allocation (x_1, \dots, x_n) where each x_i is a nonnegative rational number. The problem of computing the recovery probability for this allocation and a given subset size r is equivalent to the counting version of the following decision problem (which happens to be polynomial-time solvable):

Definition. LARGEST r -SUBSET SUM (LRSS)

Instance: Finite n -vector (a_1, \dots, a_n) with $a_i \in \mathbb{Z}_0^+$, file size $d \in \mathbb{Z}^+$, and subset size $r \in \mathbb{Z}^+$, where all a_i and d can be written as decimal numbers of length at most ℓ .

Question: Is there an r -subset $\mathbf{r} \subseteq \{1, \dots, n\}$ that satisfies $\sum_{i \in \mathbf{r}} a_i \geq d$?

Note that the allocation and file size have been scaled so that the problem parameters are all integers. To show that the counting problem #LRSS is #P-complete, we essentially apply the proof of Proposition 1, substituting #LSS with #LRSS, and stipulating that the subset size $r = m + k$ in the Turing reduction. ■

Proof of Lemma 4: Summing up the c inequalities of (18) produces

$$\sum_{j=1}^c \sum_{i \in \mathbf{r}_j} x_i \geq c.$$

The terms on the left-hand side can be regrouped to obtain

$$\sum_{i \in S} \sum_{j=1}^c \mathbb{1}[i \in \mathbf{r}_j] x_i \geq c.$$

Substituting (19) into the above inequality yields

$$\sum_{i \in S} b x_i \geq c,$$

as required. ■

Proof of Lemma 5: Let \mathcal{R} be the collection of all $\binom{n}{r}$ possible r -subsets of $\{1, \dots, n\}$. If $P_S = 1$, then any feasible allocation must satisfy

$$\sum_{i \in \mathbf{r}} x_i \geq 1 \quad \forall \mathbf{r} \in \mathcal{R}.$$

Observe that each element in $\{1, \dots, n\}$ appears the same number of times among the r -subsets in \mathcal{R} . Specifically, the number of r -subsets that contain element $i \in \{1, \dots, n\}$ is just the number of ways of choosing the other $(r-1)$ elements of the r -subset from the remaining $(n-1)$ elements of $\{1, \dots, n\}$, i.e.,

$$\sum_{\mathbf{r} \in \mathcal{R}} \mathbb{1}[i \in \mathbf{r}] = \binom{n-1}{r-1} \quad \forall i \in \{1, \dots, n\}.$$

Applying Lemma 4 with $S = \{1, \dots, n\}$, $c = \binom{n}{r}$, and $b = \binom{n-1}{r-1}$ therefore produces

$$\sum_{i=1}^n x_i \geq \frac{\binom{n}{r}}{\binom{n-1}{r-1}} = \frac{n}{r}$$

for any feasible allocation. Now, $(\frac{1}{r}, \dots, \frac{1}{r})$ is a feasible allocation since it has a recovery probability of exactly 1; because it uses the minimum possible total amount of storage $\frac{n}{r}$, this allocation is also optimal. ■

Proof of Theorem 5: Suppose that n is a multiple of r ; let positive integer α be defined such that $n = \alpha r$.

We will first prove that $P_S > 1 - \frac{r}{n}$ is a sufficient condition for the optimality of $(\frac{1}{r}, \dots, \frac{1}{r})$ by showing that if the constraint

$$\sum_{i \in \mathbf{r}} x_i \geq 1 \tag{35}$$

is satisfied for more than $(1 - \frac{r}{n}) \binom{n}{r}$ distinct r -subsets $\mathbf{r} \subseteq \{1, \dots, n\}$, then the allocation $(\frac{1}{r}, \dots, \frac{1}{r})$ minimizes the required budget T . Our approach is motivated by the observation of Lemma 4. We begin by constructing a collection of r -subsets such that if constraint (35) is satisfied for the r -subsets in this collection, then $\sum_{i=1}^n x_i \geq \frac{n}{r}$. We then demonstrate that such a collection of r -subsets can be found among any collection of more than $(1 - \frac{r}{n}) \binom{n}{r}$ distinct r -subsets.

Let

$$Q \triangleq (\mathbf{v}_1, \dots, \mathbf{v}_\alpha)$$

be an ordered partition of $\{1, \dots, n\}$ that comprises α parts, where $|\mathbf{v}_j| = r$, $j = 1, \dots, \alpha$. For a given ordered partition Q , we specify a collection of α distinct r -subsets

$$\mathcal{R}_Q \triangleq \{\mathbf{r}_1, \dots, \mathbf{r}_\alpha\},$$

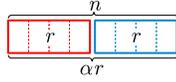
$$\text{where } \mathbf{r}_j \triangleq \mathbf{v}_j, \quad j = 1, \dots, \alpha.$$

Fig. 9 provides an example of how Q and \mathcal{R}_Q are constructed. Let A be the total number of possible ordered partitions Q . By counting the number of ways of picking \mathbf{v}_j , we have

$$A = \underbrace{\binom{\alpha r}{r} \binom{(\alpha-1)r}{r} \binom{(\alpha-2)r}{r} \dots \binom{r}{r}}_{\alpha \text{ terms}} = \frac{(\alpha r)!}{(r!)^\alpha}.$$

Let $(n, r) = (8, 4)$.

Writing $n = \alpha r$ gives $\alpha = 2$.



An example of an ordered partition is

$$Q = (\{1, 2, 3, 4\}, \{5, 6, 7, 8\}).$$

Its corresponding collection of r -subsets is

$$\mathcal{R}_Q = \{\{1, 2, 3, 4\}, \{5, 6, 7, 8\}\}.$$

Fig. 9. Example for the construction of the ordered partition Q and its corresponding collection of r -subsets \mathcal{R}_Q , in the proof of Theorem 5 (when n is a multiple of r).

Let B be the number of ordered partitions Q for which $\mathbf{r} \in \mathcal{R}_Q$, for a given r -subset $\mathbf{r} \subseteq \{1, \dots, n\}$. By counting the number of ways of picking \mathbf{v}_j , subject to the requirement that $\mathbf{r} \in \mathcal{R}_Q$, we have

$$B = \underbrace{\alpha \binom{(\alpha-1)r}{r} \binom{(\alpha-2)r}{r} \dots \binom{r}{r}}_{(\alpha-1) \text{ terms}} = \frac{\alpha((\alpha-1)r)!}{(r!)^{\alpha-1}}.$$

We claim that for any given ordered partition Q , if

$$\sum_{i \in \mathbf{r}} x_i \geq 1 \quad \forall \mathbf{r} \in \mathcal{R}_Q,$$

then $\sum_{i=1}^n x_i \geq \frac{n}{r}$. To see this, observe that each element $i \in \{1, \dots, n\}$ appears in exactly one of the α r -subsets of \mathcal{R}_Q , i.e.,

$$\sum_{\mathbf{r} \in \mathcal{R}_Q} \mathbb{1}[i \in \mathbf{r}] = 1 \quad \forall i \in \{1, \dots, n\}.$$

Applying Lemma 4 with $S = \{1, \dots, n\}$, $c = \alpha$, and $b = 1$ therefore produces $\sum_{i=1}^n x_i \geq \frac{n}{r}$.

Let \mathcal{R} be the collection of all $\binom{n}{r}$ possible r -subsets of $\{1, \dots, n\}$. Observe that all A collections \mathcal{R}_Q can be found in \mathcal{R} , i.e.,

$$\mathcal{R}_{Q_1} \subseteq \mathcal{R}, \quad \mathcal{R}_{Q_2} \subseteq \mathcal{R}, \quad \dots, \quad \mathcal{R}_{Q_A} \subseteq \mathcal{R}.$$

With each removal of an r -subset from \mathcal{R} , we reduce the number of collections \mathcal{R}_Q that can be found among the remaining r -subsets by at most B . It follows that the minimum number of r -subsets that need to be removed from \mathcal{R} so that no collections \mathcal{R}_Q remain is at least $\lceil \frac{A}{B} \rceil$, where

$$\frac{A}{B} = \frac{(\alpha r)!}{\alpha r! ((\alpha-1)r)!} = \frac{r}{n} \binom{n}{r}.$$

Thus, if fewer than $\frac{A}{B} = \frac{r}{n} \binom{n}{r}$ r -subsets are removed from \mathcal{R} , then at least one collection \mathcal{R}_Q would remain; equivalently, some collection \mathcal{R}_Q can be found among any collection of more than $(1 - \frac{r}{n}) \binom{n}{r}$ distinct r -subsets.

We have therefore shown that if $P_S > 1 - \frac{r}{n}$, then any feasible allocation must satisfy $\sum_{i=1}^n x_i \geq \frac{n}{r}$. Now, $(\frac{1}{r}, \dots, \frac{1}{r})$ is a feasible allocation since it has a recovery probability of exactly 1; because it uses the minimum possible total amount of storage $\frac{n}{r}$, this allocation is also optimal.

We proceed to prove that $P_S > 1 - \frac{r}{n}$ is also a necessary condition for the optimality of $(\frac{1}{r}, \dots, \frac{1}{r})$ by demonstrating that this allocation is suboptimal for any $P_S \leq 1 - \frac{r}{n}$.

For $r < n$, the allocation $(0, \frac{1}{r}, \dots, \frac{1}{r})$ has a recovery probability of $\binom{n-1}{r} / \binom{n}{r} = 1 - \frac{r}{n}$ and is therefore a feasible allocation for any $P_S \leq 1 - \frac{r}{n}$. Since this allocation uses a smaller total amount of storage $\frac{n-1}{r} < \frac{n}{r}$, it is a strictly better allocation than $(\frac{1}{r}, \dots, \frac{1}{r})$ for any $P_S \leq 1 - \frac{r}{n}$.

For the trivial case $r = n$, we have $1 - \frac{r}{n} = 0$. The empty allocation $(0, \dots, 0)$ is clearly optimal for any $P_S \leq 0$. ■

Proof of Theorem 6: Suppose that n is not a multiple of r ; let integers α and r' be as defined in the theorem. For brevity, we additionally define positive integers d , m , and m' such that

$$d = \gcd(r, r'), \quad r = md, \quad r' = m'd.$$

We can therefore write $n = (\alpha m + m')d$.

We will prove that

$$P_S > 1 - \frac{d}{\alpha d + m'd} = 1 - \frac{1}{\alpha + m'}$$

is a sufficient condition for the optimality of $(\frac{1}{r}, \dots, \frac{1}{r})$ by showing that if the constraint

$$\sum_{i \in \mathbf{r}} x_i \geq 1$$

is satisfied for more than $(1 - \frac{1}{\alpha+m'}) \binom{n}{r}$ distinct r -subsets $\mathbf{r} \subseteq \{1, \dots, n\}$, then the allocation $(\frac{1}{r}, \dots, \frac{1}{r})$ minimizes the required budget T . We apply the proof technique of Theorem 5, but modify the construction of the ordered partition Q and its corresponding collection of r -subsets \mathcal{R}_Q to take into account the indivisibility of n by r .

For the moment, we will proceed with the assumption that $\alpha \geq 1$. Let

$$Q \triangleq (\mathbf{u}_1, \dots, \mathbf{u}_{m'}, \mathbf{v}_1, \dots, \mathbf{v}_\alpha)$$

be an ordered partition of $\{1, \dots, n\}$ that comprises $(m' + \alpha)$ parts, where

$$\begin{aligned} |\mathbf{u}_j| &= d, & j &= 1, \dots, m', \\ |\mathbf{v}_j| &= r = md, & j &= 1, \dots, \alpha. \end{aligned}$$

For a given ordered partition Q , we specify a collection of $(m' + \alpha)$ distinct r -subsets

$$\mathcal{R}_Q \triangleq \{\mathbf{r}_1, \dots, \mathbf{r}_{m'}, \mathbf{r}_{m'+1}, \dots, \mathbf{r}_{m'+\alpha}\},$$

$$\text{where } \mathbf{r}_j \triangleq \begin{cases} \bigcup_{\ell=0}^{m-1} \mathbf{u}_{j+\ell} & \text{if } j = 1, \dots, m', \\ \mathbf{v}_{j-m'} & \text{if } j = m'+1, \dots, m'+\alpha, \end{cases}$$

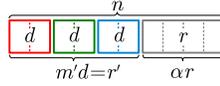
$$\text{and } \mathbf{u}_j \triangleq \mathbf{u}_{j-m'} \text{ if } j > m'.$$

Fig. 10 provides an example of how Q and \mathcal{R}_Q are constructed. Let A be the total number of possible ordered partitions Q . By counting the number of ways of picking \mathbf{u}_j and \mathbf{v}_j , we

Let $(n, r) = (10, 4)$.

Writing $n = \alpha r + r'$ gives $\alpha = 1$ and $r' = 6$.

We have $d = \gcd(r, r') = 2$, $m = r/d = 2$, and $m' = r'/d = 3$.



An example of an ordered partition is

$$Q = (\{1, 2\}, \{3, 4\}, \{5, 6\}, \{7, 8, 9, 10\}).$$

Its corresponding collection of r -subsets is

$$\mathcal{R}_Q = \left\{ \begin{array}{l} \{1, 2\}, \quad \{3, 4\}, \\ \{3, 4\}, \quad \{5, 6\}, \\ \{5, 6\}, \quad \{1, 2\}, \\ \{7, 8, 9, 10\} \end{array} \right\}.$$

Fig. 10. Example for the construction of the ordered partition Q and its corresponding collection of r -subsets \mathcal{R}_Q , in the proof of Theorem 6 (when n is not a multiple of r).

have

$$A = \underbrace{\binom{(\alpha m + m')d}{d} \binom{(\alpha m + m' - 1)d}{d} \dots \binom{(\alpha m + 1)d}{d}}_{m' \text{ terms}}.$$

$$\underbrace{\binom{\alpha m d}{m d} \binom{(\alpha - 1)m d}{m d} \dots \binom{m d}{m d}}_{\alpha \text{ terms}} = \frac{((\alpha m + m')d)!}{(d!)^{m'} ((m d)!)^{\alpha}}.$$

Let B be the number of ordered partitions Q for which $\mathbf{r} \in \mathcal{R}_Q$, for a given r -subset $\mathbf{r} \subseteq \{1, \dots, n\}$. By counting the number of ways of picking \mathbf{u}_j and \mathbf{v}_j , subject to the requirement that $\mathbf{r} \in \mathcal{R}_Q$, we have

$$B = \underbrace{\binom{((\alpha - 1)m + m')d}{d} \binom{((\alpha - 1)m + m' - 1)d}{d} \dots \binom{((\alpha - 1)m + 1)d}{d}}_{m' \text{ terms}}.$$

$$\alpha \underbrace{\binom{(\alpha - 1)m d}{m d} \binom{(\alpha - 2)m d}{m d} \dots \binom{m d}{m d}}_{(\alpha - 1) \text{ terms}}$$

$$+ m' \underbrace{\binom{m d}{d} \binom{(m - 1)d}{d} \dots \binom{d}{d}}_{m \text{ terms}}.$$

$$\underbrace{\binom{((\alpha - 1)m + m')d}{d} \binom{((\alpha - 1)m + m' - 1)d}{d} \dots \binom{(\alpha m + 1)d}{d}}_{(m' - m) \text{ terms}}.$$

$$\underbrace{\binom{\alpha m d}{m d} \binom{(\alpha - 1)m d}{m d} \dots \binom{m d}{m d}}_{\alpha \text{ terms}}$$

$$= \alpha \frac{(((\alpha - 1)m + m')d)!}{(d!)^{m'} ((m d)!)^{\alpha - 1}} + m' \frac{(((\alpha - 1)m + m')d)!}{(d!)^{m'} ((m d)!)^{\alpha - 1}}$$

$$= (\alpha + m') \frac{(((\alpha - 1)m + m')d)!}{(d!)^{m'} ((m d)!)^{\alpha - 1}}.$$

We claim that for any given ordered partition Q , if

$$\sum_{i \in \mathbf{r}} x_i \geq 1 \quad \forall \mathbf{r} \in \mathcal{R}_Q,$$

then $\sum_{i=1}^n x_i \geq \frac{n}{r}$. To see this, consider the partition of

$\{1, \dots, n\}$ formed by sets U and V , where

$$U \triangleq \bigcup_{j=1}^{m'} \mathbf{u}_j, \quad V \triangleq \bigcup_{j=1}^{\alpha} \mathbf{v}_j.$$

Correspondingly, we partition \mathcal{R}_Q into two collections of r -subsets \mathcal{R}_Q^U and \mathcal{R}_Q^V , where

$$\mathcal{R}_Q^U \triangleq \{\mathbf{r}_1, \dots, \mathbf{r}_{m'}\}, \quad \mathcal{R}_Q^V \triangleq \{\mathbf{r}_{m'+1}, \dots, \mathbf{r}_{m'+\alpha}\}.$$

Observe that each element $i \in U$ appears in exactly one \mathbf{u}_j , which in turn appears in exactly m of the m' r -subsets of \mathcal{R}_Q^U (namely $\mathbf{r}_j, \mathbf{r}_{j-1}, \dots, \mathbf{r}_{j-(m-1)}$, where $\mathbf{r}_\ell \triangleq \mathbf{r}_{\ell+m'}$ if $\ell < 1$), i.e.,

$$\sum_{\mathbf{r} \in \mathcal{R}_Q^U} \mathbb{1}[i \in \mathbf{r}] = m \quad \forall i \in U.$$

Applying Lemma 4 with $S = U$, $c = m'$, and $b = m$ therefore produces $\sum_{i \in U} x_i \geq \frac{m'}{m} = \frac{r'}{r}$. Likewise, observe that each element $i \in V$ appears in exactly one of the α r -subsets of \mathcal{R}_Q^V , i.e.,

$$\sum_{\mathbf{r} \in \mathcal{R}_Q^V} \mathbb{1}[i \in \mathbf{r}] = 1 \quad \forall i \in V.$$

Applying Lemma 4 with $S = V$, $c = \alpha$, and $b = 1$ therefore produces $\sum_{i \in V} x_i \geq \alpha$. Combining the sums of U and V yields

$$\sum_{i=1}^n x_i = \sum_{i \in U} x_i + \sum_{i \in V} x_i \geq \frac{r'}{r} + \alpha = \frac{n}{r}.$$

Let \mathcal{R} be the collection of all $\binom{n}{r}$ possible r -subsets of $\{1, \dots, n\}$. As demonstrated in the proof of Theorem 5, if fewer than $\frac{A}{B}$ r -subsets are removed from \mathcal{R} , then at least one collection \mathcal{R}_Q can be found among the remaining r -subsets. In this case, we have

$$\frac{A}{B} = \frac{1}{\alpha + m'} \frac{((\alpha m + m')d)!}{(((\alpha - 1)m + m')d)! (m d)!} = \frac{1}{\alpha + m'} \binom{n}{r}.$$

Thus, some collection \mathcal{R}_Q can be found among any collection of more than $\left(1 - \frac{1}{\alpha + m'}\right) \binom{n}{r}$ distinct r -subsets.

We have therefore shown that if $P_S > 1 - \frac{1}{\alpha + m'}$, then any feasible allocation must satisfy $\sum_{i=1}^n x_i \geq \frac{n}{r}$. Now, $(\frac{1}{r}, \dots, \frac{1}{r})$ is a feasible allocation since it has a recovery probability of exactly 1; because it uses the minimum possible total amount of storage $\frac{n}{r}$, this allocation is also optimal.

Applying the preceding argument to the degenerate case of $\alpha = 0$ produces $\frac{A}{B} = \frac{1}{m'} \binom{n}{r}$, which is consistent with the above expression. ■

Proof of Corollary 3: Suppose that n is a multiple of $(n - r)$; let integer $\beta \geq 2$ be defined such that $n = \beta(n - r) \iff n = \frac{\beta}{\beta - 1}r$.

If $\beta = 2$, then $n = 2r$, i.e., n is a multiple of r . According to Theorem 5, $(\frac{1}{r}, \dots, \frac{1}{r})$ is an optimal allocation if and only if

$$P_S > 1 - \frac{r}{n} = 1 - \frac{r}{2r} = \frac{1}{2} = \frac{r}{n},$$

as required.

If $\beta \geq 3$, then n is not a multiple of r . We can write $n = \alpha r + r'$, where $\alpha = 0$ and $r' = n \in \{r + 1, \dots, 2r - 1\}$.

According to Theorem 6, $(\frac{1}{r}, \dots, \frac{1}{r})$ is an optimal allocation if

$$P_S > 1 - \frac{\gcd(r, r')}{\alpha \gcd(r, r') + r'} = 1 - \frac{\gcd(r, n)}{n} = 1 - \frac{n-r}{n} = \frac{r}{n}.$$

To show that $P_S > \frac{r}{n}$ is also a necessary condition for the optimality of $(\frac{1}{r}, \dots, \frac{1}{r})$, we demonstrate that this allocation is suboptimal for any $P_S \leq \frac{r}{n}$. The allocation $(1, 0, \dots, 0)$ has a recovery probability of $\binom{n-1}{r-1} / \binom{n}{r} = \frac{r}{n}$ and is therefore a feasible allocation for any $P_S \leq \frac{r}{n}$. Since this allocation uses a smaller total amount of storage $1 < \frac{n}{r}$, it is a strictly better allocation than $(\frac{1}{r}, \dots, \frac{1}{r})$ for any $P_S \leq \frac{r}{n}$. ■

Proof of Lemma 6: At $T = \frac{n}{r}$, the recovery probability corresponding to a particular choice of $\ell \in \{1, 2, \dots, r-1\}$ is given by

$$P_S \left(n, r, T = \frac{n}{r}, \ell \right) = \mathbb{P} \left[\mathcal{B} \left(r, \frac{\ell}{r} \right) \geq \ell \right].$$

We will prove that the above expression is at most $\frac{3}{4}$ for any $\ell \in \{1, 2, \dots, r-1\}$ and $r \geq 2$ by showing that

$$\mathbb{P} \left[\mathcal{B} \left(a+b, \frac{a}{a+b} \right) \geq a \right] \leq \frac{3}{4}$$

for any positive integers a and b . To do this, we consider the following three exhaustive cases separately:

Case 1: Suppose that $a \geq 18$ and $b \geq 3$. We will first derive an upper bound for $\mathbb{P} \left[\mathcal{B} \left(a+b, \frac{a}{a+b} \right) \geq a \right]$ by finding separate bounds for $\mathbb{P} \left[\mathcal{B} \left(a+b, \frac{a}{a+b} \right) = a \right]$ and $\mathbb{P} \left[\mathcal{B} \left(a+b, \frac{a}{a+b} \right) \geq a+1 \right]$; we then proceed to show that this upper bound is smaller than $\frac{3}{4}$ for any $a \geq 18$ and $b \geq 3$.

For any positive integers a and b , we have

$$\begin{aligned} \mathbb{P} \left[\mathcal{B} \left(a+b, \frac{a}{a+b} \right) = a \right] &= \binom{a+b}{a} \left(\frac{a}{a+b} \right)^a \left(\frac{b}{a+b} \right)^b \\ &< \frac{e^{\frac{1}{12(a+b)}}}{\sqrt{2\pi}} \sqrt{\frac{a+b}{ab}}. \end{aligned} \quad (36)$$

Inequality (36) follows from the application of the following bound for the binomial coefficient:

$$\binom{a+b}{a} < \frac{e^{\frac{1}{12(a+b)}}}{\sqrt{2\pi}} \frac{(a+b)^{a+b+\frac{1}{2}}}{a^{a+\frac{1}{2}} b^{b+\frac{1}{2}}},$$

which is derived from the following Stirling-based bounds for the factorial (see, e.g., [31]):

$$\sqrt{2\pi k} \left(\frac{k}{e} \right)^k < k! < \sqrt{2\pi k} \left(\frac{k}{e} \right)^k e^{\frac{1}{12k}}, \quad k \geq 1.$$

For any positive integers a and b , we have

$$\mathbb{P} \left[\mathcal{B} \left(a+b, \frac{a}{a+b} \right) \geq a+1 \right] \leq \frac{1}{2}, \quad (37)$$

which follows from the definition of the median: The mean of the binomial random variable $\mathcal{B} \left(a+b, \frac{a}{a+b} \right)$ is $(a+b) \cdot \frac{a}{a+b} = a$; since the mean is an integer, the median coincides with the mean [32]. Therefore, according to the definition of the median, we have

$$\mathbb{P} \left[\mathcal{B} \left(a+b, \frac{a}{a+b} \right) \leq a \right] \geq \frac{1}{2},$$

which leads to inequality (37).

Combining bounds (36) and (37) produces

$$\mathbb{P} \left[\mathcal{B} \left(a+b, \frac{a}{a+b} \right) \geq a \right] < \frac{e^{\frac{1}{12(a+b)}}}{\sqrt{2\pi}} \sqrt{\frac{a+b}{ab}} + \frac{1}{2} \triangleq f(a, b)$$

for any positive integers a and b . Now, the upper bound $f(a, b)$ is a decreasing function of both a and b since $f(a, b)$ is a symmetric function and the partial derivative

$$\frac{\partial}{\partial a} f(a, b) = -\frac{6b^2 + 6ab + a}{12a(a+b)^2} \frac{e^{\frac{1}{12(a+b)}}}{\sqrt{2\pi}} \sqrt{\frac{a+b}{ab}}$$

is negative for any $a \geq 1$ and $b \geq 1$. Thus, for any $a \geq 18$ and $b \geq 3$, we have

$$f(a, b) \leq f(a=18, b=3) = \frac{e^{\frac{1}{252}}}{6} \sqrt{\frac{7}{\pi}} + \frac{1}{2} \approx 0.749773 < \frac{3}{4},$$

which implies that $\mathbb{P} \left[\mathcal{B} \left(a+b, \frac{a}{a+b} \right) \geq a \right] < \frac{3}{4}$ for any positive integers $a \geq 18$ and $b \geq 3$.

Case 2: Suppose that $b \in \{1, 2\}$. We will show that

$$\mathbb{P} \left[\mathcal{B} \left(a+1, \frac{a}{a+1} \right) \geq a \right] \leq \frac{3}{4} \text{ and } \mathbb{P} \left[\mathcal{B} \left(a+2, \frac{a}{a+2} \right) \geq a \right] < \frac{3}{4}$$

for any positive integer a . The left-hand side of each inequality can be expanded and simplified to obtain the following:

$$\mathbb{P} \left[\mathcal{B} \left(a+1, \frac{a}{a+1} \right) \geq a \right] = \frac{a^a (2a+1)}{(a+1)^{a+1}} \triangleq f_1(a),$$

$$\mathbb{P} \left[\mathcal{B} \left(a+2, \frac{a}{a+2} \right) \geq a \right] = \frac{a^a (5a^2 + 10a + 4)}{(a+2)^{a+2}} \triangleq f_2(a).$$

The first derivatives of $f_1(a)$ and $f_2(a)$, which are given by

$$f_1'(a) = \frac{a^a}{(a+1)^{a+1}} \left\{ 2 - (2a+1) \ln \left(\frac{a+1}{a} \right) \right\},$$

$$f_2'(a) = \frac{a^a}{(a+2)^{a+2}} \left\{ (10a+10) - (5a^2 + 10a + 4) \ln \left(\frac{a+2}{a} \right) \right\},$$

can be shown to be negative for any $a \geq 1$. Since $f_1(a=1) = \frac{3}{4}$, $f_2(a=1) = \frac{19}{27} < \frac{3}{4}$, and both $f_1(a)$ and $f_2(a)$ are decreasing functions of a for any $a \geq 1$, it follows that $f_1(a) \leq \frac{3}{4}$ and $f_2(a) < \frac{3}{4}$ for any positive integer a , as required.

Case 3: Suppose that $a \in \{1, 2, \dots, 17\}$. We will describe our approach for $a=1$ and $a=2$; the proofs for the other 15 cases are similar, and can be verified with the help of a computer. We will show that

$$\mathbb{P} \left[\mathcal{B} \left(b+1, \frac{1}{b+1} \right) \geq 1 \right] \leq \frac{3}{4} \text{ and } \mathbb{P} \left[\mathcal{B} \left(b+2, \frac{2}{b+2} \right) \geq 2 \right] < \frac{3}{4}$$

for any positive integer b . The left-hand side of each inequality can be expanded and simplified to obtain the following:

$$\mathbb{P} \left[\mathcal{B} \left(b+1, \frac{1}{b+1} \right) \geq 1 \right] = 1 - \frac{b^{b+1}}{(b+1)^{b+1}} \triangleq g_1(b),$$

$$\mathbb{P} \left[\mathcal{B} \left(b+2, \frac{2}{b+2} \right) \geq 2 \right] = 1 - \frac{b^{b+1} (3b+4)}{(b+2)^{b+2}} \triangleq g_2(b).$$

The first derivatives of $g_1(b)$ and $g_2(b)$, which are given by

$$g_1'(b) = \frac{b^b}{(b+1)^{b+1}} \left\{ b \ln \left(\frac{b+1}{b} \right) - 1 \right\},$$

$$g_2'(b) = \frac{b^b}{(b+2)^{b+2}} \left\{ (3b^2 + 4b) \ln \left(\frac{b+2}{b} \right) - (6b+4) \right\},$$

can be shown to be negative for any $b \geq 1$. Since $g_1(b=1) = \frac{3}{4}$, $g_2(b=1) = \frac{20}{27} < \frac{3}{4}$, and both $g_1(b)$ and $g_2(b)$ are decreasing functions of b for any $b \geq 1$, it follows that

$g_1(b) \leq \frac{3}{4}$ and $g_2(b) < \frac{3}{4}$ for any positive integer b , as required. ■

Proof of Theorem 7: We have already established that the choice of $\ell = r$ is optimal for any $T \geq \frac{n}{r}$; it therefore suffices to show that $\ell = r$ is also optimal for any $T \in \left[\frac{n}{r} \left(\frac{3}{4} \right)^{\frac{1}{r}}, \frac{n}{r} \right)$.

The recovery probability corresponding to any $\ell \in \{1, 2, \dots, r\}$ is given by

$$P_S(n, r, T, \ell) = \mathbb{P} \left[\mathcal{B} \left(r, \min \left(\frac{\ell T}{n}, 1 \right) \right) \geq \ell \right],$$

which is a nondecreasing function of T since $\min \left(\frac{\ell T}{n}, 1 \right)$ either increases or remains constant at 1 as T increases. More precisely, $P_S(n, r, T, \ell)$ is an increasing function of T on the interval $(0, \frac{n}{\ell})$; for higher values of T , the function saturates at 1. We can verify this claim by checking that the partial derivative

$$\frac{\partial}{\partial p} \mathbb{P} [\mathcal{B}(r, p) \geq \ell] = \ell \binom{r}{\ell} p^{\ell-1} (1-p)^{r-\ell}$$

is positive for any $p \in (0, 1)$.

Now, the recovery probability corresponding to the choice of $\ell = r$ at $T = \frac{n}{r} \left(\frac{3}{4} \right)^{\frac{1}{r}}$ is given by

$$P_S \left(n, r, T = \frac{n}{r} \left(\frac{3}{4} \right)^{\frac{1}{r}}, \ell = r \right) = \mathbb{P} \left[\mathcal{B} \left(r, \left(\frac{3}{4} \right)^{\frac{1}{r}} \right) \geq r \right] = \frac{3}{4}.$$

Since $P_S(n, r, T, \ell)$ is a nondecreasing function of T , we have

$$P_S(n, r, T, \ell = r) \geq \frac{3}{4} \quad \text{for any } T \geq \frac{n}{r} \left(\frac{3}{4} \right)^{\frac{1}{r}}.$$

On the other hand, for any $\ell \in \{1, 2, \dots, r-1\}$, we have

$$P_S(n, r, T, \ell) \leq \frac{3}{4} \quad \text{for any } T \leq \frac{n}{r},$$

from the upper bound of Lemma 6. It therefore follows that the choice of $\ell = r$ is optimal for any $T \in \left[\frac{n}{r} \left(\frac{3}{4} \right)^{\frac{1}{r}}, \frac{n}{r} \right)$, as required. ■

Proof of Corollary 4: Theorem 7 already demonstrates that the choice of $\ell = r$ is optimal for any $T \geq \frac{n}{r} \left(\frac{3}{4} \right)^{\frac{1}{r}}$; we will proceed to show that a recovery probability of at least $\frac{3}{4}$ is *not* achievable for any $T < \frac{n}{r} \left(\frac{3}{4} \right)^{\frac{1}{r}}$.

Recall from the proof of Theorem 7 that the recovery probability $P_S(n, r, T, \ell)$ corresponding to any $\ell \in \{1, 2, \dots, r\}$ is an increasing function of T on the interval $(0, \frac{n}{\ell})$. Thus, for the choice of $\ell = r$, the function $P_S(n, r, T, \ell = r)$ increases wrt T on the subinterval $\left(0, \frac{n}{r} \left(\frac{3}{4} \right)^{\frac{1}{r}} \right) \subset \left(0, \frac{n}{r} \right)$; since $P_S \left(n, r, T = \frac{n}{r} \left(\frac{3}{4} \right)^{\frac{1}{r}}, \ell = r \right) = \frac{3}{4}$, it follows that

$$P_S(n, r, T, \ell = r) < \frac{3}{4} \quad \text{for any } T < \frac{n}{r} \left(\frac{3}{4} \right)^{\frac{1}{r}}.$$

On the other hand, for any $\ell \in \{1, 2, \dots, r-1\}$, the function $P_S(n, r, T, \ell)$ increases wrt T on the subinterval $\left(0, \frac{n}{\ell} \right) \subset \left(0, \frac{n}{r} \right)$; since $P_S \left(n, r, T = \frac{n}{r}, \ell \right) \leq \frac{3}{4}$ according to Lemma 6, it follows that

$$P_S(n, r, T, \ell) < \frac{3}{4} \quad \text{for any } T < \frac{n}{r}.$$

Hence, the optimal recovery probability for any $T < \frac{n}{r} \left(\frac{3}{4} \right)^{\frac{1}{r}}$ is strictly less than $\frac{3}{4}$. ■

ACKNOWLEDGMENT

The authors would like to thank Brighten Godfrey and Robert Kleinberg for introducing the problem to them and for sharing their insights. The authors also thank Dimitris Achlioptas for the interesting discussions.

REFERENCES

- [1] D. Leong, A. G. Dimakis, and T. Ho, "Distributed storage allocations," *IEEE Trans. Inf. Theory*, vol. 58, no. 7, pp. 4733–4752, Jul. 2012.
- [2] —, "Distributed storage allocation problems," in *Proc. Workshop Netw. Coding, Theory, and Appl. (NetCod)*, Jun. 2009.
- [3] —, "Distributed storage allocation for high reliability," in *Proc. IEEE Int. Conf. Commun. (ICC)*, May 2010.
- [4] —, "Symmetric allocations for distributed storage," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, Dec. 2010.
- [5] A. G. Dimakis, P. B. Godfrey, Y. Wu, M. J. Wainwright, and K. Ramchandran, "Network coding for distributed storage systems," *IEEE Trans. Inf. Theory*, vol. 56, no. 9, pp. 4539–4551, Sep. 2010.
- [6] A. Jiang, "Network coding for joint storage and transmission with minimum cost," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jul. 2006.
- [7] T. Ho, M. Médard, R. Koetter, D. R. Karger, M. Effros, J. Shi, and B. Leong, "A random linear network coding approach to multicast," *IEEE Trans. Inf. Theory*, vol. 52, no. 10, pp. 4413–4430, Oct. 2006.
- [8] C. Fragouli, J.-Y. L. Boudec, and J. Widmer, "Network coding: An instant primer," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 1, pp. 63–68, Jan. 2006.
- [9] R. Kleinberg, R. Karp, C. Papadimitriou, and E. Friedman, *Personal correspondence between R. Kleinberg and A. G. Dimakis*, Oct. 2006.
- [10] J.-S. Wu and R.-J. Chen, "An algorithm for computing the reliability of weighted- k -out-of- n systems," *IEEE Trans. Reliability*, vol. 43, no. 2, pp. 327–328, Jun. 1994.
- [11] Y. Chen and Q. Yang, "Reliability of two-stage weighted- k -out-of- n systems with components in common," *IEEE Trans. Reliability*, vol. 54, no. 3, pp. 431–440, Sep. 2005.
- [12] W. K. Lin, D. M. Chiu, and Y. B. Lee, "Erasure code replication revisited," in *Proc. Int. Conf. Peer-to-Peer Comput. (P2P)*, Sep. 2004.
- [13] A. Tsigros and Z. J. Haas, "Analysis of multipath routing—Part I: The effect on the packet delivery ratio," *IEEE Trans. Wireless Commun.*, vol. 3, no. 1, pp. 138–146, Jan. 2004.
- [14] —, "Analysis of multipath routing, Part 2: Mitigation of the effects of frequently changing network topologies," *IEEE Trans. Wireless Commun.*, vol. 3, no. 2, pp. 500–511, Mar. 2004.
- [15] S. Jain, M. Demmer, R. Patra, and K. Fall, "Using redundancy to cope with failures in a delay tolerant network," in *Proc. ACM SIGCOMM*, Aug. 2005.
- [16] M. Sardari, R. Restrepo, F. Fekri, and E. Soljanin, "Memory allocation in distributed storage networks," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2010.
- [17] N. Alon, P. Frankl, H. Huang, V. Rödl, A. Ruciński, and B. Sudakov, "Large matchings in uniform hypergraphs and the conjectures of Erdős and Samuels," *arXiv:1107.1219*, Jul. 2011. [Online]. Available: <http://arxiv.org/abs/1107.1219>
- [18] P. Erdős, "A problem on independent r -tuples," *Ann. Univ. Sci. Budapest*, vol. 8, pp. 93–95, 1965.
- [19] M. Naor and R. M. Roth, "Optimal file sharing in distributed networks," *SIAM J. Comput.*, vol. 24, no. 1, pp. 158–183, Feb. 1995.
- [20] A. Jiang and J. Bruck, "Network file storage with graceful performance degradation," *ACM Trans. Storage*, vol. 1, no. 2, pp. 171–189, May 2005.
- [21] D. Leong, T. Ho, and R. Cathey, "Optimal content delivery with network coding," in *Proc. Annu. Conf. Inf. Sci. Syst. (CISS)*, Mar. 2009.
- [22] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Spray and Wait: An efficient routing scheme for intermittently connected mobile networks," in *Proc. ACM SIGCOMM Workshop Delay-Tolerant Netw.*, Aug. 2005.
- [23] S. Acedánski, S. Deb, M. Médard, and R. Koetter, "How good is random linear coding based distributed networked storage?" in *Proc. Workshop Netw. Coding, Theory, and Appl. (NetCod)*, Apr. 2005.
- [24] A. G. Dimakis, V. Prabhakaran, and K. Ramchandran, "Ubiquitous access to distributed data in large-scale sensor networks through decentralized erasure codes," in *Proc. Int. Symp. Inf. Process. Sensor Netw. (IPSN)*, Apr. 2005.
- [25] A. Kamra, V. Misra, J. Feldman, and D. Rubenstein, "Growth codes: Maximizing sensor network data persistence," in *Proc. ACM SIGCOMM*, Sep. 2006.

- [26] Y. Lin, B. Liang, and B. Li, "Data persistence in large-scale sensor networks with decentralized fountain codes," in *Proc. IEEE Int. Conf. Comput. Commun. (INFOCOM)*, May 2007.
- [27] S. A. Aly, Z. Kong, and E. Soljanin, "Fountain codes based distributed storage algorithms for large-scale wireless sensor networks," in *Proc. ACM/IEEE Int. Conf. Inf. Process. Sensor Netw. (IPSN)*, Apr. 2008.
- [28] L. G. Valiant, "The complexity of computing the permanent," *Theoretical Comput. Sci.*, vol. 8, no. 2, pp. 189–201, 1979.
- [29] M. Sipser, *Introduction to the Theory of Computation*, 2nd ed. Boston, MA: Thomson Course Technology, 2006.
- [30] M. Mitzenmacher and E. Upfal, *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. New York: Cambridge Univ. Press, 2005.
- [31] W. Feller, *An Introduction to Probability Theory and Its Applications, Vol. 1*, 3rd ed. New York: Wiley, 1968.
- [32] R. Kaas and J. M. Buhrman, "Mean, median and mode in binomial distributions," *Statistica Neerlandica*, vol. 34, no. 1, pp. 13–18, 1980.