# D-Rex: Heterogeneity-Aware Reliability Framework and Adaptive Algorithms for Distributed Storage

## Maxime Gonthier

University of Chicago Chicago, IL, USA Argonne National Laboratory Lemont, IL, USA mgonthier@uchicago.edu

### **Bogdan Nicolae**

Argonne National Laboratory Lemont, IL, USA bogdan.nicolae@acm.org

## Valerie Hayot-Sasson

University of Chicago Chicago, IL, USA Argonne National Laboratory Lemont, IL, USA vhayot@uchicago.edu

## Dante D.

Sanchez-Gallegos University Carlos III of Madrid Leganes, Spain dantsanc@pa.uc3m.es

## Sicheng Zhou

Southern University of Science and Technology Shenzhen, China zhousc2021@mail.sustech.edu.cn

## Greg Pauloski

University of Chicago Chicago, IL, USA jgpauloski@uchicago.edu

## Kyle Chard

University of Chicago Chicago, IL, USA Argonne National Laboratory Lemont, IL, USA chard@uchicago.edu

## Abstract

The exponential growth of data necessitates distributed storage models, such as peer-to-peer systems and data federations. While distributed storage can reduce costs and increase reliability, the heterogeneity in storage capacity, I/O performance, and failure rates of storage resources makes their efficient use a challenge. Further, node failures are common and can lead to data unavailability and even data loss. Erasure

## 

This work is licensed under a Creative Commons Attribution 4.0 International License. *ICS '25, Salt Lake City, UT, USA* © 2025 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-1537-2/25/06 https://doi.org/10.1145/3721145.3730412

## Haochen Pan

University of Chicago Chicago, IL, USA haochenpan@uchicago.edu

## Hai Duc Nguyen

University of Chicago Chicago, IL, USA Argonne National Laboratory Lemont, IL, USA ndhai@uchicago.edu

## Jesus Carretero

University Carlos III of Madrid Leganes, Spain jcarrete@inf.uc3m.es

## Ian Foster

University of Chicago Chicago, IL, USA Argonne National Laboratory Lemont, IL, USA foster@anl.gov

coding is a common resiliency strategy implemented in storage systems to mitigate failures by striping data across storage locations. However, erasure coding is computationally expensive and existing systems do not consider the heterogeneous resources and their varied capacity and performance when placing data chunks. We tackle the challenges of using erasure coding with distributed and heterogeneous nodes, aiming to store as much data as possible, minimize encoding and decoding time, and meeting user-defined reliability requirements for each data item. We propose two new dynamic scheduling algorithms, D-Rex LB and D-Rex SC, that adaptively choose erasure coding parameters and map chunks to heterogeneous nodes. D-Rex SC achieves robust performance for both storage utilization and throughput, at a higher computational cost, while D-Rex LB is faster but with slightly less competitive performance. In addition, we propose two greedy algorithms, GreedyMinStorage and GreedyLeastUsed, that

optimize for storage utilization and load balancing, respectively. Our experimental evaluation shows that our dynamic schedulers store, on average, 45% more data items without significantly degrading I/O throughput compared to state-ofthe-art algorithms, while GreedyLeastUsed is able to store 21% more data items while also increasing throughput.

#### **CCS** Concepts

Computer systems organization → Reliability; • Hardware → External storage; • Information systems → Distributed storage; • Theory of computation → Scheduling algorithms.

#### Keywords

Fault Tolerance, Reliability, Distributed & Heterogeneous Data Storage, Erasure Coding, Load Balancing

#### **ACM Reference Format:**

Maxime Gonthier, Dante D. Sanchez-Gallegos, Haochen Pan, Bogdan Nicolae, Sicheng Zhou, Hai Duc Nguyen, Valerie Hayot-Sasson, Greg Pauloski, Jesus Carretero, Kyle Chard, and Ian Foster. 2025. D-Rex: Heterogeneity-Aware Reliability Framework and Adaptive Algorithms for Distributed Storage. In 2025 International Conference on Supercomputing (ICS '25), June 08–11, 2025, Salt Lake City, UT, USA. ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/ 3721145.3730412

#### 1 Introduction

The increasing size of data, coupled with its distributed production and acquisition, is driving the adoption of distributed data storage models. Distributed storage is often seen as a way of reducing costs, increasing fault tolerance, and improving performance. However, achieving these benefits requires addressing the challenges posed by heterogeneous storage nodes (potentially geographically distributed) with varying capacities, bandwidths, and failure rates.

A key challenge in this context is the resilient storage of data in a cost-efficient manner. Erasure coding (EC) [23] is a well-known technique that provides the same resilience guarantees as full data replication but uses significantly less storage space at the cost of additional computations. EC techniques divide a data item *d* into *K* chunks of size size(d)/K and add *P* parity chunks of the same size, such that any *K* of the K + P chunks are sufficient to recover the original item through a decoding operation. Thus, if each of the K + P chunks is placed on a different storage system, up to *P* failures can be tolerated before data items are lost.

However, EC faces two main challenges. First, it incurs costs due to the complexity of chunking objects and calculating the parity data [24, 28]: costs that tend to increase with the number of chunks. To illustrate these costs, we conducted an experiment in which we encoded and decoded



Figure 1: Breakdown of encode, decode, and upload times to store 400 MB across a varying number of nodes based on [28] on a 48-core Intel Xeon E5-2670. The parity chunks P = 2, resulting in K data chunks equal to total number of chunks minus two.

a 400 MB data item with a fixed P = 2 and varying values of K. As shown in Figure 1, increased K reduces storage overhead but increases decoding time. Second, state-of-theart distributed storage algorithms employ EC using fixed values of K and P and distribute chunks uniformly, under the assumption that all data items are equal and the storage nodes are homogeneous. Given the diversity of storage requirements (e.g., some data items require a higher reliability guarantee and/or accessed faster than others) and heterogeneity of storage nodes (e.g., due to federation of storage systems), this assumption is outdated and necessitates more adaptive EC.

To address these challenges, we propose D-Rex (Dynamic Resilience Extension), an adaptive EC algorithm based on an innovative reliability model that allows a customized reliability target for each stored data item. Based on this reliability model, we contribute two greedy algorithms and two adaptive algorithms that automatically decide the number of data chunks, parity chunks, and allocation of chunks on storage nodes in order to satisfy reliability targets, maximize storage space utilization, and minimize the I/O overhead of data access. D-Rex is grounded in real-world applications and can be integrated into existing data transfer and storage infrastructures, such as Globus [2] or DynoStore [35, 36]. We summarize our contributions as follows:

- (1) We introduce an innovative reliability model that leverages the notion of a *reliability target*, expressed as the chance of successful access to a data item within a given time interval, making it easy for users to express and reason about quality-of-service guarantees (§3.1).
- (2) We introduce two greedy algorithms, GreedyMinStorage and GreedyLeastUsed, and two dynamic algorithms, D-Rex LB and D-Rex SC, based on our reliability model. The

D-Rex algorithms aim to solve the multi-objective optimization problem (§3.2) of successfully storing as much data as possible, satisfying the reliability target and storage capacity constraints, and minimizing I/O overheads associated with read/write operations (§4).

- (3) We design and implement a dynamic data storage simulator that we employ to evaluate our proposed algorithms against the state-of-the-art algorithms (§5.2) under controlled settings. Using the simulator, we run extensive simulations using diverse datasets, showing that D-Rex SC and D-Rex LB store on average 45% and 31% more data while only 0.4 and 0.3 MB/s slower, respectively, than state-of-the-art algorithms (§5) when using heterogeneous nodes. GreedyLeastUsed increases the amount of data stored by up to 21% while also increasing throughput.
- (4) We integrate our algorithms into a real memory system to perform experiments in an unconstrained memory scenario and show that our algorithms do not significantly degrade throughput under such conditions (§6).

#### 2 Related Work

Distributed Storage: Distributed object stores and file systems often rely on replication, where full copies of a data item are stored on different nodes to ensure availability. AWS S3 [1] replicates objects across at least three availability zones. Similarly, Hadoop Distributed File System (HDFS) replicates each data item twice [6], resulting in a total of three copies and a storage overhead of 200%. The Inter Planetary File System [11] is a peer-to-peer (P2P) file system that fully replicates data items as they are shared between different peers and, by the nature of P2P, must manage many heterogeneous nodes. SkyStore [27] dynamically replicates and evicts objects across different cloud storage providers according to application access patterns with the goal of optimizing for lower cost. Our solutions differ by using erasure coding, which has lower storage overhead than full replication.

**Erasure Coding in Distributed Storage:** HDFS [3] uses the well-established Reed-Solomon [41] erasure coding scheme and stores data and parity chunks in a distributed fashion to minimize inter-rack write traffic. It relies on a centralized NameNode to manage the encoding process. Gluster [33] is a distributed file system that combines disk storage from multiple servers into a single namespace. Gluster integrates erasure coding with a default configuration of 4 data chunks and 2 parity chunks. It uses dispersed bricks rather than dedicated nodes to store chunks, allowing for more granular storage. Distributed Asynchronous Object Storage (DAOS) is an object store where objects are managed as collections, called containers, that can scale across multiple storage devices [17]. It implements EC and allows users to

define the desired level of reliability. The storage efficiency of EC has motivated its implementation in distributed data storage systems such as EdgeHydra [16] and ELECT [34]. In these systems, erasure codes stripe data items across multiple storage nodes to balance workloads and enable parallel data access for improved performance. OceanStore [21] and ER-Store [26] implement a combination of EC and replication to provide high availability of data, assuming that all nodes are prone to fail. While these systems efficiently manage EC, they assume that all nodes are homogeneous, and as a result, the number of data and parity chunks can be static. Our approach differs by incorporating heterogeneity into decisions on how to utilize EC and where to place chunks.

**Erasure Coding Optimizations:** SMFRepair [42] suggests saturating the bandwidth of the least busy node and transferring data from links with the largest possible bandwidth to minimize the transfer time when decoding data. We consider a simpler model where we do not include transfers between storage nodes. [19] propose a linear programming model for the data block placement problem at the edge in order to minimize storage cost while maintaining quality of service. Some approaches use TOPSIS to optimize chunk placement on heterogeneous storage nodes to improve read/write throughput [32] but do not consider storage size limitations.

While these solutions optimize for either transfer time or storage cost, we have developed an approach that optimizes for both time and space in the context of data storage on heterogeneous nodes with limited storage.

**Failure-Recovery Techniques** Erasure coding (EC) significantly lowers repair bandwidth compared to full replication [40]. To further reduce recovery latency, schemes such as proactive repair rebuild lost chunks immediately upon failure detection, or even preemptively when failures are predicted [25, 37]. These techniques focus solely on the repair phase, and can be layered on top of our chunking and placement model without any modifications. Checkpoint-restart is another actively researched area. Several efforts leverage multiple storage levels to asynchronously hide the I/O overhead of persisting checkpoints [15, 29, 30]. EC is used to protect the node-local checkpoints to a parallel file system.

#### 3 Framework

Here, we formalize the problem statement. First, we introduce an innovative reliability model that addresses the question of satisfying resilience based on user-specified quality of service constraints. Then, we formulate an optimization problem to decide how to encode and distribute multiple data items that need to be written to a resilient data repository

Notation	Description					
Known to the algorithms						
$L \\ S = \{S_1, \dots, S_L\}$ size(S <sub>i</sub> ) F(S <sub>i</sub> , t)	Number of available storage nodes. The set of storage nodes. Total storage size of node $S_i$ . Free storage size on node $S_i$ at time					
$Pr_{failure}(S_i, \Delta t)$	Probability of node $S_i$ failing at least once over a time period $\Delta t$ .					
$B_r(S_i) B_w(S_i) T_{encode}(N_d, K_d, size(d))$	Read bandwidth of $S_i$ , in MB/s. Write bandwidth of $S_i$ , in MB/s. Time required to encode a data item <i>d</i> of size <i>size(d)</i> into $N_d$ chunks,					
$T_{\mathbf{decode}}(K_d, size(d))$	each of size $size(d)/K_d$ . Time required to recreate a data item of size $size(d)$ from $K_d$ chunks.					
Known for each data i	tem $d$ at the time that it is to be stored					
size(d) $\Delta t_d$ RT(d)	Size of $d$ in MB. Duration $d$ has to stay available. Reliability target of $d$ .					
Unknown to the algor	ithms					
$m$ $\mathbb{D} = \{d_1, \dots, d_m\}$	Number of data items to store. The set of data items to store.					
Chosen by the algorith	nms for each data item d					
K <sub>d</sub> P <sub>d</sub>	Number of data chunks required to retrieve data item. Controls the data item's chunk size (= $\lceil size(d)/K_d \rceil$ ). Number of parity chunks added for					
$egin{array}{cc} C_d & & \ N_d & & \ \mathcal{M}_d & & \ \end{array}$	<i>d</i> . Set of chunks for data item <i>d</i> . Total number of chunks: $N_d = K_d + P_d$ . Subset of nodes in S where the chunks of $C_d$ have been loaded. $ C_d  =  \mathcal{M}_d  = K_d + P_d$					

Table 1: List of notations.

subject to our reliability model. Key notations are summarized in Table 1.

#### 3.1 Reliability Model

As a first contribution, we introduce a reliability model that allows the user to input the level of reliability they want for their data. The goal is to give the user as much flexibility as possible without over-complicating how the user must reason about resilience. Intuitively, each data item that needs to be stored has a well-defined life cycle. For example, the intermediate results of an application need to persist for the lifetime of the application. Experimental results from scientific applications typically need to be persistent for days or months until data are analyzed and summarized. Some important results need to be archived and retained for years.



Figure 2: Data storage and retrieval, from top left, clockwise, with algorithm decisions in blue. (1) Encoding of data d into N = K data chunks + P parity chunks; (2) Writing of chunks to storage nodes, with choices shown by  $\rightarrow$ ; (3) Reading of K chunks from selected nodes (more choices,  $\leftarrow$ ); and (4) Decoding of original content from read chunks.

Thus, each data item *d* (out of a total of *m* data items) to be stored in the repository has a retention time  $\Delta t_d < \infty$  (i.e., an expiration). We propose the notion of a reliability target RT(d), a value between (0, 1) representing the chance that *d* can be successfully retrieved before expiration. The user only specifies  $\Delta t_d$  and RT(d).

A repository consists of *L* storage nodes  $S = \{S_1, \ldots, S_L\}$ . A storage node may fail at any time. We assume an erasurecoded representation that splits each data item d into a chunking  $C_d$  of  $K_d$  equally sized data chunks, complemented by  $P_d$  parity chunks of the same size. Each chunk is stored on a different node. We denote the mapping of chunks to nodes as  $\mathcal{M}_d$ . If at least  $K_d$  chunks of the total  $K_d + P_d$  chunks are still accessible on nodes before d expires, the data item is available. For example, Figure 2 shows the encoding of a data item into 3 data chunks and 2 parity chunks that are written to 5 distinct nodes with mapping  $(S_1, S_3, S_5, S_7, S_9)$ . Assuming  $S_5$  has failed, the data item can be reconstructed from the two remaining data blocks and any parity block using a decoding operation. Encoding and decoding involve overheads  $T_{encode}$  and  $T_{decode}$ , respectively. Note that the erasure-coded representation is the most general one. For example, replication can be seen as a particular case in which  $K_d = 1$  represents the original data item,  $P_d$  is the number of replicas, and  $T_{encode} = T_{decode} = 0$ .

We assume a fail-stop failure model in which the storage nodes are either operational or have permanently failed. Failure results in the permanent loss of any chunks stored on the node. For simplicity, the probability of node failure is equivalent to the probability of disk failure (detailed in §5.3). We assume the remaining components in the node will not fail, but this model can be extended to account for such failures. Lastly, we assume that the failure rate of a storage node is constant. Then, the probability of a storage node failing at least once during a period  $\Delta t$  (expressed as a fraction of a year):

$$Pr_{\text{failure}}(S_i, \Delta t) = 1 - e^{-\lambda_{\text{rate}} \cdot \Delta t}$$
(1)

Assuming failures follow a homogeneous Poisson process with a constant rate of occurrence  $\lambda_{rate}$  [13], we can compute the probability of a data item d with  $P_d$  parity chunks and a mapping  $\mathcal{M}_d$  to remain available for  $\Delta t_d$  days:  $Pr_{avail}(\mathcal{M}_d, P_d, \Delta t_d)$ . This is equivalent to the probability of no more than  $P_d$  nodes failing. To this end, we use the cumulative Poisson binomial distribution function to calculate the probability of a certain number X of events (i.e., node failures in our case) across a set of independent Bernoulli trials with probability  $Pr_{failure}(S_d, \Delta t_d)$ :

$$Pr_{\text{avail}}(\mathcal{M}_d, P_d, \Delta t_d) = Pr(X \le P_d) = \sum_{l=0}^{P_d} \sum_{A \in F_l} \prod_{i \in A} Pr_{\text{failure}}(S_i, \Delta t_d) \prod_{j \in A^c} (1 - Pr_{\text{failure}}(S_j, \Delta t_d))$$

$$(2)$$

with  $F_l$  the set of all possible subsets of nodes in  $\mathcal{M}_d$  containing exactly l nodes. Equation 2 is precise but quadratic in complexity, so we implemented an approximation inspired by previous work [18, 38]. This approximation yields the *reliability constraint* of each data item, defined as:

$$Pr_{\text{avail}}(\mathcal{M}_d, P_d, \Delta t_d) \ge RT(d)$$
 (3)

Equation 3 is a key optimization constraint based on a simple user-specified reliability target.

While we assume uniform and independent node failures, our framework extends to any failure-prediction model (e.g., as in [8]). We would only need to replace  $Pr_{\text{failure}}(S_i, \Delta t)$  with the predicted probability of node  $S_i$  to fail, given its current age, within the next  $\Delta t$ . Equation 2 could then be applied without any modification.

#### 3.2 **Problem Formulation**

We assume the *m* requests to store data items are ordered. Multiple users may submit requests concurrently, but for the purposes of space allocation, each item must be considered sequentially. This does not preclude physically writing many items concurrently. Each data item in  $\mathbb{D} = \{d_1, \ldots, d_m\}$  is associated with a submission timestamp  $t_i > t_j$  for each i > j, repository-assigned based on concurrency control policies.

For each storage node, the following information is available: its total capacity,  $size(S_i)$ ; free space at time t,  $F(S_i, t)$ ; write bandwidth,  $B_w(S_i)$ ; read bandwidth,  $B_r(S_i)$ ; and probability of failure over a given  $\Delta t$ ,  $Pr_{failure}(S_i, \Delta t)$ : see Equation 1. Our goal is to automatically determine for each data item d: (i) the number of data chunks,  $K_d$ ; (ii) the number of parity chunks,  $P_d$ ; and (iii) the mapping  $\mathcal{M}_d$  of these chunks onto a subset of  $K_d + P_d$  storage nodes in  $\mathbb{S}$ , while optimizing two quality metrics: (i) the total size of data **items success-fully stored** and (ii) the **I/O overhead** of the write/read operations. Furthermore, the optimization must be done in an online fashion, meaning any decision about  $d_i$  can only depend on decisions taken for previous data items  $d_j$  with j < i, i.e., foreknowledge of storage requests is not available.

The motivation behind the quality metrics is based on two observations: (i) the goal of a storage repository is to maximize effective use of its capacity and (ii) users desire requests to be served quickly in addition to successfully. These two quality metrics may require a trade-off: increasing the success rate of write requests may come at the cost of greater I/O overhead. We do not prioritize one metric because both are important.

More formally, we consider a write operation successful if and only if it simultaneously satisfies two constraints: (i) there is capacity available on each node in  $\mathcal{M}_d$  to store the new chunk and (ii) the reliability constraint  $Pr_{\text{avail}}(\mathcal{M}_d, P_d,$  $\Delta t_d) \geq RT(d)$  (introduced in Equation 3) can be satisfied. Then, the total amount of data successfully stored is defined as:

 $\mathbb{W} = \sum_{d \in \mathcal{D}_{\text{succass}}} size(d)$ 

where:

$$\mathcal{D}_{\text{success}} = d \in \mathbb{D} \mid \operatorname{and} \forall S_i \in \mathcal{M}_d, \ F(S_i, t_d) \ge \frac{size(d)}{K_d}$$

We are concerned with I/O overheads  $\mathbb{T}$  required to encode, decode, write, and read each successfully stored data item. These overheads depend on  $K_d$ ,  $P_d$  and  $\mathcal{M}_d$ . The relative importance of each overhead can be weighed depending on the application and user requirements, but we assume overheads share equal importance for the purposes of this work. Thus, we define average I/O throughput as  $\mathbb{T}$  where:

$$\mathbb{T} = \frac{\mathbb{W}}{\sum_{d \in \mathcal{D}_{\text{success}}} T_{\text{encode}} + T_{\text{decode}} + T_{\text{write}} + T_{\text{read}}}$$

We assume that all chunk transfers are parallelized and that no two nodes share a connection, which would introduce contention. Thus, the slowest node in  $\mathcal{M}_d$  is the bottleneck

for both writing ( $T_{write}$ ) and reading ( $T_{read}$ ). We assume that encoding and decoding times are constant for a given configuration  $K_d$ ,  $P_d$ , and size(d).

Given the formal definitions of  $\mathbb{W}$  and  $\mathbb{T}$ , we summarize the optimization problem in 1.

PROBLEM 1. Given a sequence of m data items  $\mathbb{D} = \{d_1, \ldots, d_m\}$  and L storage nodes  $\mathbb{S} = \{S_1, \ldots, S_L\}$ , determine a **chunk**ing  $C_d$  of each data item  $d \in \mathbb{D}$  as a set of  $K_d$  data blocks and  $P_d$  parity blocks, plus a **mapping**  $\mathcal{M}_d$  that assigns each data block and parity block to a storage node, such as to maximize the average I/O throughput  $\mathbb{T}$  and the total amount  $\mathbb{W}$  of successfully stored data items while subject to capacity availability on each storage node in  $\mathcal{M}_d$  to store the chunks in  $C_d$  and a target reliability constraint defined as  $\mathcal{M}_d$  satisfy  $Pr_{avail}(\mathcal{M}_d, P_d, \Delta t_d) \geq RT(d)$ .

#### 4 Dynamic Resilience Extension (D-Rex)

To address the limitations described in §2, in this section we present four algorithms that aim to solve the problem described in §3.2. The first two algorithms we propose are greedy algorithms (§4.1 and §4.2). To address the limitations of these greedy algorithms, we propose two more algorithms. The first algorithm, presented in §4.3, aims to achieve good load balancing while having low complexity, making it easily applicable. The second algorithm, presented in §4.4, is more computationally expensive but seeks to optimize the tradeoff between throughput and amount of data stored. For each algorithm, we describe below the decision made when a new data item *d* is stored at time *t*.

#### 4.1 Greedy Storage Minimization

We assume that a reasonable solution for a cost minimization problem like Problem 1 is to minimize the total storage overhead for each data item. We denote this algorithm *GreedyMinStorage*. It selects values for  $N_d$  and  $K_d$  that minimize the storage overhead for each data item by solving an optimization problem such as:

minimize 
$$\frac{size(d)}{K_d} \times N_d$$
  
s.t.  $Pr_{\text{avail}}(\mathcal{M}_d, K_d, \Delta t_d) \ge RT(d)$  (4)

and use  $P_d = N_d - K_d$  for the number of parity chunks. The mapping  $\mathcal{M}_d$  is chosen by favoring the nodes with the fastest bandwidth. This usually corresponds to selecting high values of N and K, so as to make each chunk as small as possible. The intuition is that, first, smaller chunks are faster to transfer, so in the multi-objective problem we defined, this would ensure read/write efficiency. Second, it incurs minimal storage overhead, allowing for more data items to be stored, thus satisfying the storage objective we defined.

#### 4.2 Greedy Maximum Free Space

The *GreedyLeastUsed* algorithm prioritizes nodes with the most available storage when placing chunks. This approach ensures that data is distributed to the least utilized nodes, balancing storage consumption across the system, which we believe is a reasonable approach that a system manager would want to use. The algorithm selects values for  $P_d$  and  $K_d$  that are as small as possible in order to greedily reduce the encoding/decoding overhead:

minimize 
$$P_d + K_d$$
  
s.t.  $Pr_{\text{avail}}(\mathcal{M}_d, K_d, \Delta t_d) \ge RT(d),$  (5)  
and  $\mathcal{M}_d = K_d + P_d$  nodes with highest  $F(S_i)$ 

#### 4.3 D-Rex LB: Adding Load Balancing

The greedy algorithms presented above do not account for varying storage sizes. The intuition behind D-Rex LB is to calculate a balance penalty for all nodes, regardless of whether a node stores a given item. By assigning penalties to nodes that do not store the data item, the algorithm promotes a more even distribution of load across the system, enhancing load balancing and overall system efficiency.

The first step of the algorithm, detailed in 1, is to compute the average free capacity ( $F_{\text{average}}$ ) of the storage nodes. We then iterate over the number of data chunks  $K_d$  required to reconstruct the original data item. For nodes that would store the data item d, we calculate a balance penalty based on the node's free storage, the average free storage across all nodes, and the chunk size. For nodes that would not store d, we still compute the penalty but exclude the chunk size. When the average free capacity is high, indicating that most nodes are relatively empty, the chunk size has minimal impact on the balance penalty, and the algorithm prioritizes filling the largest nodes first. Conversely, when the average free storage is small, the nodes begin to saturate and the chunk size has a greater impact on the balance penalty. This incentives the algorithm to use larger values of  $K_d$ . Ultimately, we choose the value of  $K_d$  that yields the smallest penalty while also minimizing  $P_d$  and ensuring the reliability target is met.

D-Rex LB requires looping  $\frac{L(L+1)}{2}$  times over each mapping. Verifying whether a mapping satisfies the reliability constraint requires computing the CDF for the Poisson-Binomial distribution, which in our implementation has a cost of at most  $O(K^2)$ . Since K can be at most L - 1, the worst-case complexity simplifies to  $O(L^4)$ . However, in most cases K is small and the break condition at Line 23 is reached early. Consequently, the scheduling overhead per data item is negligible compared to the data transfer time, as summarized in Table 2. D-Rex: Heterogeneity-Aware Reliability Framework and Adaptive Algorithms for Distributed Storage

ICS '25, June 08-11, 2025, Salt Lake City, UT, USA

Algorithm 1 D-Rex Load Balancing (D-Rex LB)

1:  $F_{\text{average}} \leftarrow \frac{1}{L} \sum_{i=1}^{L} F(S_i, t_d)$ 2: Sort  $\mathbb{S}$  in order of decreasing  $F(S_i, t_d)$ 3: min  $bp \leftarrow \infty$ // bp: balance penalty 4:  $min_K \leftarrow -1$ 5: **for**  $P_d \leftarrow 1$  to  $P_d < L$  **do for**  $K_d \leftarrow 2$  to  $K_d \leq L - P_d$  **do** 6: 7:  $bp \leftarrow 0$  $\mathcal{M}_d \leftarrow \text{first } K_d + P_d \text{ nodes in sorted } \mathbb{S}$ 8: **if**  $Pr_{avail}(\mathcal{M}_d, P_d, \Delta t_d) \geq RT(d)$  **then** 9: for each  $S_i \in \mathcal{M}_d$  do 10:  $bp += |F(S_i, t_d) - size(d)/K_d - F_{average}|$ 11: end for 12: **for** each  $S_i \notin \mathcal{M}_d$  **do** 13:  $bp += \left| F(S_j, t_d) - F_{\text{average}} \right|$ 14: end for 15: **if** *bp* < *min\_bp* **then** 16:  $min_K \leftarrow K_d$ 17:  $min\_bp \leftarrow bp$ 18: end if 19: end if 20: end for 21: if min  $K \neq -1$  then 22: Break 23: end if 24: 25: end for 26: Create  $min_K + P_d$  chunks of size  $size(d)/min_K$  and store them on the first  $min_K + P_d$  nodes of sorted S

4.4 D-Rex SC: Considering System Capacity

We propose a second dynamic algorithm, *D-Rex SC* (see 2), that is more computationally intensive than D-Rex LB but adapts better to node storage saturation. The intuition is to consider three key parameters: (i) the cost of encoding and decoding, (ii) the erasure coding storage overhead for a given N and K, and (iii) the proportion of free storage remaining across the nodes.

We begin by considering different node mappings (e.g., [1, 2, 3], [1, 2, 4], [1, 2, 3, 4], etc.) as described in Lines 2–8. We consider the first 2<sup>10</sup> mappings of nodes with the highest available storage, starting with the top nodes sequentially. As an example, initially, we examine mappings of nodes [1, 2], then expand to [1, 2, 3] and so on up to [1, 2, ..., L]. Then, we explore other combinations such as [2, 3] etc. Limiting to 2<sup>10</sup> mappings ensures a balance between complexity and effective results at scale.

For each mapping  $\mathcal{M}$ , we know the number of nodes,  $N = |\mathcal{M}|$ , and the failure rate of each node in  $\mathcal{M}$ . We then calculate the couple K, P that minimize storage overhead

while meeting the data reliability goal as well as the three key parameters: duration, storage cost, and saturation cost.

Duration comprises the encoding and decoding times and the parallelized read and write times (assumed to be the time for the node with the worst bandwidth) using the nodes in  $\mathcal{M}$ . Linear regression is used to estimate the encode and decode times given a set of encode and decode times from different data item sizes and numbers of data and parity chunks. In practice, our tests over data sizes ranging from 1 to 500 MB show that the linear regression prediction closely matches the actual measurements. Storage cost is the storage overhead of erasure coding for N nodes with K data chunks. Saturation cost aims to evaluate how close a node is to its storage limit. We assign a saturation score ranging from 0 to 1 based on the node's utilization relative to its maximum capacity. The scoring function (denoted as f(x) on line 8) follows an exponential curve to penalize nodes approaching their limit (see Figure 3). Completely filling a node is often a bad solution, as it reduces the number of nodes that can be used for future data items, thus forcing the use of smaller values of N, which is worse overall for storage overhead.

These parameters are interrelated, and optimizing one can lead to trade-offs in the others. For instance, reducing storage overhead with larger values of N and K increases the decoding time. Therefore, instead of choosing solutions that excel in only one parameter, we retain only those on the Pareto front (Line 12) that we call candidate mappings. In Line 16, we compute for all candidates the progress that they made on each parameter relative to the other candidates. Thus, if all candidates have a similar or equal storage overhead, they will not make much progress relative to others. This favors solutions that are radically better than others on one or more parameters.

Finally, a score is calculated that combines the progress of the three parameters. In order to account for the available capacity of the nodes, we compute the total system saturation (see line 11) using the same exponential function as before. However, instead of evaluating individual nodes, we compute the overall free space across all storage nodes. This system saturation is used to diminish the importance given to the duration parameter. As a result, this penalizes solutions that try to use parameters that are costly in terms of storage when the system is already saturated. The candidate solution with the highest score is selected.

Choosing K (Line 4) requires iterating over L - 1 possible values of K and computing the CDF for the Poisson-Binomial distribution. This results in a complexity of  $O(L^3)$  for this step. This step is repeated for each mapping, which is at most 2<sup>10</sup> times. This does not affect the asymptotic complexity, but it does affect the scheduling overhead for smaller values of L. Adding the loop Line 15, the overall complexity is  $O(max(L^3, |\mathbb{C}|))$ . However, the scheduling overhead per

Algorithm 2 D-Rex System Capacity (D-Rex SC)

- 0: Task: Determine for supplied d the  $\mathcal{M}, K, P$  that provide best trade-off between amount of data stored and throughput
- 1: Sort S in order of decreasing  $F(S_i, t_d)$
- $\mathbb{M}_{set} \leftarrow \text{First } 2^{10} \text{ combinations of nodes}$ 2:
- for each mapping  $\mathcal{M} \in \mathbb{M}_{set}$  do 3:
- Choose (K, P) such that  $\lceil \frac{size(d)}{K} \rceil \times (K+P)$  is minimized and  $Pr_{avail}(\mathcal{M}, P, \Delta t) \ge RT(d)$ 4:
- $C \leftarrow (K, P, \mathcal{M})$ 5:
- $duration(C) \leftarrow min(B_{w} \in \mathcal{M}) \times \lceil \frac{size(d)}{K} \rceil + min(B_{r} \in \mathcal{M}) \times \lceil \frac{size(d)}{K} \rceil // Compute \ duration \ for \ given \ d, \ \mathcal{M}, K + T_{encode}(|\mathcal{M}|, K, size(d)) + T_{decode}(K, size(D))$ 6:
- 7:

 $storage(C) \leftarrow \lceil \frac{size(d)}{K} \rceil \times |\mathcal{M}|$   $saturation(C) \leftarrow \sum_{k=1}^{|\mathcal{M}|} f(size(S_k) - F(S_k, t) + \lceil \frac{size(d)}{K} \rceil)$ Add C to candidate\_mappings 8:

9:

// Compute storage for given d, M, K

// Compute saturation for given d, M, K

- 10: end for
- System Saturation  $\leftarrow f(size(\mathbb{S}) F(\mathbb{S}))$  for  $f(x) = \exp(\min(size(d) \in \mathbb{D} \text{ at time } t), 1/L)(size(\mathbb{S}), 1)$ 11:
- $\mathbb{C} \leftarrow$  Pareto front of all  $C \in candidate\_mappings$  based on {duration(C), storage(C), saturation(C)} 12:
- $(\text{duration}_{\min}, \text{ storage}_{\min}, \text{ saturation}_{\min}) \leftarrow (\min_{c \in \mathbb{C}} duration(c), \min_{c \in \mathbb{C}} storage(c), \min_{c \in \mathbb{C}} saturation(c))$ 13:
- $(\text{duration}_{max}, \text{ storage}_{max}, \text{ saturation}_{max}) \leftarrow (\max_{c \in \mathbb{C}} duration(c), \max_{c \in \mathbb{C}} storage(c), \max_{c \in \mathbb{C}} saturation(c))$ 14:
- 15: **for** each candidate mapping  $C_i \in \mathbb{C}$  **do**
- $duration(C_i) duration_{min}$  $(duration_{\text{progress}}(C_i), storage_{\text{progress}}(C_i), saturation_{\text{progress}}(C_i))$ (1 16: 1  $\frac{\text{storage}_{(C_i)-\text{storage}_{min}}{\text{storage}_{max}-\text{storage}_{min}}, \ 1 - \frac{\text{saturation}_{(C_i)-\text{saturation}_{min}}{\text{saturation}_{max}-\text{saturation}_{min}})$
- $\frac{1}{\text{storage}_{max} \text{storage}_{min}}, 1 \frac{1}{\text{saturation}_{max} \text{saturation}_{min}}, Score(C_i) \leftarrow (1 System\_Saturation) \times duration_{\text{progress}}(C_i) + (storage_{\text{progress}}(C_i) + saturation_{\text{progress}}(C_i))/2$ 17:

18: end for

- 19: Select candidate mapping  $C_{best} = (K_{best}, P_{best}, \mathcal{M}_{best})$  with highest score
- 20: Partition *d* into  $K_{best}$  chunks of size  $size(d)/K_{best}$ , add  $P_{best}$  parity chunks and distribute across nodes in  $\mathcal{M}_{best}$



Figure 3: Example of the function used to calculate the saturation cost. The exponential curve spans from the smallest known data item size to the total storage capacity of the node. The green dot is an example of a node's available free storage.

data item remains reasonable even as the number of nodes scales, as shown in Table 2.

#### **Experimental evaluation** 5

We evaluate our algorithms using a simulator written in C, since processing hundreds of terabytes of data in real-world

Table 2: Measured scheduling overhead per data item (in milliseconds) for varying numbers of nodes.

Alg. $\downarrow$ # of nodes $\rightarrow$	10	50	100	500
GreedyMinStorage	0.007	0.146	0.943	9.464
GreedyLeastUsed	0.001	0.003	0.006	0.027
D-Rex LB	0.001	0.015	0.057	2.462
D-Rex SC	0.488	2.871	6.671	343.3

experiments is costly. The simulator also allows us to easily explore different scenarios and node failures. We obtain data workloads (data items, sizes, arrival times) from four real-world datasets. The simulator processes data items using their release date from the dataset. It calculates transfer times using user-reported bandwidths without interference, i.e. in a best-case scenario. We implemented all strategies and state-of-the-art (SOTA) baselines in this simulator.<sup>1</sup> We calibrated the simulator with real measurements, and as will be shown by Figure 8 and Figure 13, the results measured in the simulation are confirmed by real experiments. We first

<sup>&</sup>lt;sup>1</sup>Our simulator, algorithm implementations, and details of storage nodes and datasets are available for reproducibility: https://github.com/Double-Blind-975/Drex-repro-anonymous.git.

Table 3: Number of data items and	l per-item statistics
for the four datasets used in our e	xperiments.

		Per-item statistics					
Dataset	# of items (m)	Avg.	Min	Max	Std		
MEVA	4157	117.1 MB	1.4 MB	856.1 MB	68.1 MB		
Sentinel-2	256,351	475.9 MB	2.7 MB	969.9 MB	256.5 MB		
SWIM	5214	23.4 GB	1.0 B	5329.5 GB	177.0 GB		
IBM COS	47,529	2.6 GB	0.2 MB	1345.8 GB	18.9 GB		

A]	lgorit	hm 3	Erasure-	Cod	ling
----	--------	------	----------	-----	------

- 1: **Function** EC(*K*, *P*, *d*, *t*)
- 2: Sort S by descending value of  $B_w$
- 3:  $N \leftarrow K + P$
- 4: Create *N* chunks, each of size  $\lceil \frac{size(d)}{K} \rceil$ 5: Store chunks on first *N* nodes that satisfy  $\forall S_i \in \mathcal{M}$ ,  $\lceil \frac{size(d)}{K} \rceil \leq F(S_i, t) \text{ and } Pr_{avail}(\mathcal{M}, P, \Delta t_d) \geq RT(d)$ 6: End function

describe the datasets and storage nodes used in the experiments and then evaluate our algorithms across different sets of nodes, datasets, and node failure scenarios.

#### 5.1 Datasets

We use four datasets, summarized in Table 3. The MEVA dataset [7] contains 4, 157 five-minute video clips, totaling 346 hours, captured from multiple cameras. It is widely used for human activity recognition research [10]. The Sentinel-2 dataset [12] provides multispectral Earth imagery. Imagery data are emitted and processed nearly daily, with a stable size. We sampled a representative portion of the dataset for benchmarking. The Statistical Workload Injector for MapReduce (SWIM) [5] is a tool for evaluating HDFS performance with realistic workloads, widely used in big data systems research [22]. The IBM Cloud Object Storage (IBM COS) dataset [14] captures public data such as PDF documents, media files, database backups, or disk images. It has been used to study caching policies [39]. For benchmarking consistency, we standardize the total size of all requests performed for each dataset to 122 TB, trimming the last three datasets and repeating the MEVA dataset to meet this size.

#### State-of-the-Art Algorithm Baselines 5.2

Within the scope of the proposed reliability model and problem formulation, we compare our algorithms against two state-of-the-art baseline algorithms that are widely used in modern file systems and cloud platforms: three static erasure coding (EC) approaches and another with a varied replication factor.

5.2.1 Static Erasure Coding. Popular solutions like HDFS [3] and Gluster [33] use static erasure coding to store data items. We implement the default configurations of HDFS and Gluster, utilizing 10 or fewer storage nodes, as this will be the standard number of nodes in our experimental evaluation. For HDFS, there are two erasure coding configurations: three data and two parity blocks, referred to as EC(3, 2); and six data and three parity blocks, referred to as EC(6, 3). Gluster uses an EC(4, 2) configuration. We describe how we apply this erasure coding in 3. For fairness, we assume that HDFS and Gluster prioritize nodes with the highest bandwidth.

5.2.2 Erasure Coding with Varied Replication Factors. DAOS provides users with a selection of predefined erasure coding configurations [9, 17], including EC(8, 1), EC(8, 2), EC(4, 1), and EC(4, 2). For scenarios requiring higher reliability, DAOS also supports full data replication with configurations of  $2\times$ , 4×, or 6× full copies. To ensure a fair evaluation, we prioritize choosing the configuration that meets the predefined reliability target  $Pr_{\text{avail}}(\mathcal{M}_d, P_d, \Delta t_d) \ge RT(d)$  with the lowest storage overhead, aligning their performance with our objectives.

#### 5.3 Storage Nodes from Backblaze

We use statistics from Backblaze [4] to identify commonly used storage HDDs and their failure rates (SSDs and HDDs have similar failure rates [31]). We construct four distinct storage node sets where each node contains a single drive and each set contains 10 nodes: (i) Most Used: 10 most used HDDs from Backblaze, representing a realistic scenario with popular hardware; (ii) Most Unreliable: 10 HDDs with the highest failure rates, modeling a worst-case pathological scenario; (iii) Most Reliable: 10 HDDs with the fewest failures; and (iv) Homogeneous: 10 identical HDDs, each representing the most used HDD model from Backblaze.

Figure 4 shows the distribution of storage size and annual failure rates for the Most Used and Most Unreliable sets. Storage sizes vary from 5 to 20 TB, and failure rates show significant variability for Most Unreliable. Write bandwidths vary between 100 and 250 MB/s, and read bandwidths between 100 and 400 MB/s. Table 4 shows the pairwise correlations across all HDDs. The storage properties are largely independent, displaying no strong correlations, with the exception of read and write bandwidths. Thus, it is unreasonable to simply rely on the most performant drives, which motivates exploiting the heterogeneity of the drives/nodes.

#### 5.4 Varying Reliability Target

Figure 5 shows the proportion of data stored using the Most Used nodes and MEVA dataset while varying the target reliability. At higher reliability targets, some bars are missing as some algorithms cannot deliver the required reliability



Figure 4: Distributions of storage size and annual failure rate for the *Most Used* (top) and *Most Unreliable* (bottom) sets.

Table 4: Pairwise Pearson correlation among storageproperties for all drives recorded by Backblaze.

	Storage Size	Write BW	Read BW	Failure Rate
Storage Size	1			
Write BW	0.614	1		
Read BW	0.495	0.915	1	
Failure Rate	0.078	-0.044	-0.092	1

level. In all cases, D-Rex SC and GreedyMinStorage store the most data. GreedyLeastUsed stores 10% to 30% less data items than D-Rex SC. GreedyLeastUsed spreads the load evenly but always tries to create as few data chunks as possible. As a result, larger chunks of data are created to satisfy the reliability constraint, which increases the overall storage overhead of the stored data and thus allows fewer data items to be stored overall. The SOTA algorithms (EC(3,2), EC(4,2), EC(6,4) and DAOS) perform poorly due to the fixed storage overhead associated with their fixed choice of *P* and *K*. Additionally, they do not load balance causing certain nodes to become saturated. For example, Figure 6 shows the storage distribution of EC(3,2) with a 90% reliability target. The fastest nodes (1-5) are saturated; however, as EC(3,2) requires at least five nodes, no remaining items can be stored once node 9 is full, although significant storage remains available. In contrast, D-Rex SC, D-Rex LB, and GreedyMinStorage avoid this limitation and fully utilize nodes.

Figure 5 shows that D-Rex SC effectively incorporates the strengths of a greedy approach, storing at least 73% more data than the SOTA strategies and as much data as a greedy algorithm that minimizes storage overhead.

#### 5.5 Varying Storage Node Characteristics

Figure 7 shows the amount of data stored on different node sets when each data item is assigned a reliability target (a "number of nines") chosen at random, as follows. Let *x* be a discrete random variable over the range  $\{-1, 0, 1, 2, 3, 4, 5\}$ . Define:

$$f(x) = \begin{cases} 90 & \text{if } x = -1, \\ 100 - 10^{-x} & \text{if } 0 \le x < 5, \\ 99.99999 & \text{if } x = 5. \end{cases}$$

Then, if  $x \neq 5$ , the reliability target is a random number chosen uniformly from the range [f(x), f(x+1)], else it is 99.99999.

The storage proportion results are consistent with the previous results: both D-Rex SC and GreedyMinStorage store significantly more data, even when dealing with unreliable or homogeneous nodes, than the other algorithms.

Figure 8 compares the average data storage throughput of the D-Rex algorithms with the greedy and SOTA algorithms. Since different algorithms store varying amounts of data and use different subsets of nodes, a fair comparison requires matching the amount of data stored. For instance, a load-balancing strategy may utilize slower nodes, reducing throughput compared to a strategy that stores less data but uses only the same fast storage nodes. Thus, we compare the average throughput for each algorithm for the same data items. For example, if EC(3,2) stores only the first 40% of the workload, we compare its throughput with that of D-Rex SC/LB on the same set of data items. We observe that the only algorithms faster than the D-Rex algorithms are GreedyLeastUsed, EC(3,2), and EC(4,2). The largest difference occurs on the Homogeneous node set, where algorithms using fewer nodes, such as GreedyLeastUsed and EC(3,2), benefit from reduced encoding/decoding time without the cost of not optimizing data placement for heterogeneity. This result is to be expected, as D-Rex SC/LB were not designed for homogeneous scenarios. In contrast, with diverse nodes like the Most Unreliable or Most Used, the difference with the D-Rex algorithms is negligible, approximately a 0.4 MB/s decrease at most.

To better understand the performance of the different algorithms, we show in Figure 9 a breakdown of the time spent on each data operation for a subset of the MEVA workload. (We use a subset so that all strategies can store all data items. Otherwise, a fair comparison would not be possible, as some strategies store more data and thus spend more time on data operations.) We see that time spent reading & writing is negligible compared to that spent encoding & decoding because encode & decode operations cannot be parallelized. While GreedyMinStorage saves time on reading & writing by creating smaller chunks, D-Rex LB and SC have faster encoding & decoding because they balance storage and time overhead.



D-Rex: Heterogeneity-Aware Reliability Framework and Adaptive Algorithms for Distributed Storage ICS '25, June 08–11, 2025, Salt Lake City, UT, USA

Figure 5: Evolution of proportion of data stored (the sum of the sizes of data stored by an algorithm relative to the sum of all data sizes in the workload) across varying reliability targets with the *Most Used* nodes and MEVA dataset. Missing bars indicate that the required reliability level cannot be achieved.



Figure 6: Consumed versus available storage for each node using EC(3,2) with a 90% reliability target.



Figure 7: Proportion of data stored (the sum of the sizes of data stored by an algorithm relative to the sum of all data sizes in the workload) with different sets of nodes and random reliability targets between 90% and 99.99999% assigned to each item from the MEVA dataset.

Overall, the D-Rex algorithms consistently store more data items across a heterogeneous set of nodes

than SOTA algorithms while achieving comparable or better throughput (slow-down ranging from 0.2 to 0.8 MB/s for D-Rex SC and from 0.2 to 0.4 MB/s for D-Rex LB compared to the best SOTA result). In a homogeneous scenario, although D-Rex stores more data items, it suffers greater loss in terms of throughput as our algorithms are not designed for such scenarios. In that case, GreedyLeastUsed is the better algorithm to use. Compared to a greedy algorithm that focuses on minimizing storage space, D-Rex SC stores as much data with significantly better throughput.

#### 5.6 Different Datasets

Figure 10 shows the amount of data stored using D-Rex and baseline algorithms for three different datasets: Sentinel-2, SWIM, and IBM COS (see §5.1). Figure 11 shows the difference in throughput between the D-Rex algorithms and the SOTA algorithms with the same datasets. Across several datasets with different data item counts and sizes, the D-Rex algorithms (SC and LB) store on average at least 45% and 31% more data than EC(3,2), EC(4,2), EC(6,3), or DAOS, respectively, while improving throughput or showing negligible degradation. Compared to a greedy algorithm that optimizes for storage overhead, D-Rex SC stores almost as many data items while achieving higher throughput, by an average of 1.5 MB/s across the three workloads. GreedyLeastUsed stores on average 21% more data items while improving throughput compared to the SOTA.

#### 5.7 Resilience against Node Failures

Tables 12a and 12b present the results of an experiment where we simulate the failures of varying numbers of nodes.



Figure 8: Average throughput difference when comparing D-Rex SC (plain) and D-Rex LB (hashed) against the rest of the algorithms using different sets of nodes and random reliability targets between 90% and 99.99999% assigned to each item from the MEVA. We compare equivalent data volumes by taking the minimum amount stored between the D-Rex and the compared algorithms. Values above 0 indicate that D-Rex algorithms are faster, while values below 0 indicate slower performance.

The experiment uses the *Most Unreliable* set of nodes and the MEVA dataset, which consists of 70 days of input data. To simulate node failures, we calculate the daily failure rate of each node based on its known annual failure rate. Jobs are submitted as usual, and after one day of job submissions (tracked using the known jobs' submission times), we simulate node failures. For each node, we generate a random number between 0 and 1; if the random number is less than or equal to the node's daily failure rate, the node is considered to have failed. The algorithms schedule data items as in previous experiments. When a node fails, they reschedule the lost chunks to meet the reliability target. If a data item cannot meet the target, its chunks are removed from all nodes, reducing the percentage of stored data shown in



Figure 9: Time span of the data operations with a dataset that does not saturate the nodes and a reliability target of 99.99%.



Figure 10: Proportion of data stored with the Sentinel-2, SWIM, and IBM COS datasets, a workload saturating all nodes, using the *Most Used* nodes and random reliability between 90% and 99.99999% assigned to each item.

the tables. Figure 12a shows the amount of data retained after failures with a target reliability of 90%. The two D-Rex algorithms and GreedyMinStorage retain the most data by using more storage nodes to reduce storage overhead and larger values of K and P to survive more failures. Overall, D-Rex SC retains the most data items.

When increasing the target reliability to 99.999%, as shown in Figure 12b, fewer failures can be survived across all algorithms, yet GreedyMinStorage and D-Rex SC still retain the most data items. With a fixed *K* and *P*, the static algorithms cannot achieve sufficient reliability after two nodes fail—maintaining 100% reliability until that point. Node failures change the reliability that can be achieved with a given set of nodes. Therefore, dynamic strategies that can increase the number of parity chunks and the total number of nodes used are able to retain more data items after failures. D-Rex: Heterogeneity-Aware Reliability Framework and Adaptive Algorithms for Distributed Storage ICS '25, June 08-11, 2025, Salt Lake City, UT, USA



Figure 11: Average throughput difference when comparing D-Rex SC (plain) and D-Rex LB (hashed) against the rest of the algorithms using different datasets, the *Most Used* set of nodes and random reliability targets between 90% and 99.99999% assigned to each data item.

Table 5: Infrastructure used to evaluate D-Rex on theChameleon Cloud.

Node(s)	Site	# CPUs	RAM (GB)	Storage		
Houe(s)	Site	# CI U3	Killin (OD)	Size (GB)	Drive	
0, 1	TACC	40	125	370	INTEL SSDSC1BG40	
2	TACC	40	126	2000	Seagate ST2000NX0273	
3	TACC	160	251	450	Micron MTFDDAK480TDS	
4, 5	NRP	48	125	200	Seagate ST9250610NS	
6	UC	96	251	960	Dell Express Flash CD5	
7	UC	96	251	7600	INTEL SSDPF2KX076TZ	
8	UC	96	187	240	Dell MZ7KM240HMHQ0D3	
9	UC	96	251	865	INTEL SSDPF2KX076TZ	

#### 6 Validation on Real Infrastructure

To show the effectiveness of the D-Rex algorithms, we conduct real experiments with a wide-area storage system called DynoStore (DS) [35, 36]. DynoStore implements a federated storage overlay across distributed storage nodes using common interfaces called data containers. DynoStore can be configured to use different algorithms to map data items to storage nodes. Thus, while DynoStore manages the access to data and storage resources, we use D-Rex to control the chunking and placement of data items. We deploy DynoStore with D-Rex using ten storage nodes distributed across three Chameleon Cloud sites: CHI@UC, CHI@TACC, and

(a) remainly target of 70%.								
Alg. $\downarrow$ # failure $\rightarrow$	2	3	4	5	6	7		
D-Rex SC	100%	100%	100%	56%	38%	28%		
D-Rex LB	100%	100%	100%	54%	38%	23%		
GreedyMinStorage	100%	100%	100%	52%	34%	28%		
GreedyLeastUsed	100%	100%	100%	15%	0%	0%		
EC(3,2)	100%	100%	100%	22%	0%	0%		
EC(4,2)	100%	100%	82%	0%	0%	0%		
EC(6,3)	0%	0%	0%	0%	0%	0%		
DAOS	100%	100%	100%	0%	0%	0%		
D-Rex LB GreedyMinStorage GreedyLeastUsed EC(3,2) EC(4,2) EC(6,3) DAOS	100% 100% 100% 100% 100% 100%	100% 100% 100% 100% 100% 0% 100%	100% 100% 100% 100% 82% 0% 100%	54% 52% 15% 22% 0% 0% 0%	38% 34% 0% 0% 0% 0%	23% 28% 0% 0% 0% 0%		

(a) Reliability target of 90%

(	b)	Relia	ability	target	of	99.999%
	v,	nem	to mit y	unger	UI.	//./////.

Algo $\downarrow$ # failure $\rightarrow$	2	3	4	5	6	7
D-Rex SC	100%	100%	100%	15%	0%	0%
D-Rex LB	100%	100%	100%	0%	0%	0%
GreedyMinStorage	100%	100%	100%	15%	0%	0%
GreedyLeastUsed	100%	100%	100%	15%	0%	0%
EC(3,2)	0%	0%	0%	0%	0%	0%
EC(4,2)	0%	0%	0%	0%	0%	0%
EC(6,3)	0%	0%	0%	0%	0%	0%
DAOS	85%	66%	41%	0%	0%	0%

# Figure 12: Proportion of data items retained after node failures under different reliability targets.



# Figure 13: Average throughput observed for uploading & encoding and downloading & decoding data in real infrastructure. DS means DynoStore. Averaged over 10 iterations.

CHI@NRP [20]. To achieve a heterogeneous setup, we selected nodes of various types and purposes. The hardware details are presented in Table 5.

We compared the performance of D-Rex with HDFS. For HDFS, we evaluated two erasure code policies using Reed-Solomon with parameters EC(3,2) and EC(6,3). Figure 13 shows the average throughput observed when uploading (top) and downloading (bottom) 1000 randomly generated data items with a resilience target of 99.999%. The upload operations include the encoding of objects, whereas the download operations include the decoding of objects. The total

size of the dataset is 113.5 GB with an average data item size of 116.3 MB. D-Rex SC/LB, GreedyLeastUsed, and HDFS with EC(3,2) achieved similar performance: in Figure 13, for data uploading, these algorithms yield a throughput of 42.4, 41.3, 41.9, and 44.2 MB/s, respectively. HDFS with EC(3,2) performs slightly better because D-Rex utilizes a larger number of nodes, including some that are slower. Since the workload does not saturate the nodes, the advantages observed in simulations, where D-Rex achieved superior performance, are not yet apparent. When increasing the number of nodes used with HDFS EC(6,3), we observe that HDFS performs worse than D-Rex. In particular, D-Rex SC and D-Rex LB are 21% and 19% faster, respectively. A similar trend is observed in download operations. D-Rex stores more data items, retains more data items after failures (as demonstrated in the simulation results), and produces similar throughput to static algorithms from the literature (as demonstrated via experiments).

#### 7 Conclusion

Distributed storage systems use erasure coding (EC) to increase reliability and reduce costs. However, EC is computationally intensive, and the varying capacity, performance, and failure rates of heterogeneous nodes present placement challenges. Existing solutions rarely take node heterogeneity into account when using EC. We address this gap by proposing two new dynamic schedulers that jointly select EC parameters and assign chunks to nodes: D-Rex-SC, which optimizes for overall system capacity, and D-Rex-LB, which balances load across nodes. Through rigorous simulations, we showed that D-Rex is able to store on average 45% more data items across different datasets and sets of nodes compared to classic state-of-the-art algorithms. Through real-world experiments, we showed that D-Rex achieves comparable throughput compared to HDFS's erasure coding. D-Rex SC provides the best overall performance compared to classic state-of-the-art algorithms in terms of both storage utilization and throughput, although its complexity increases with the number of nodes. In contrast, D-Rex LB offers a balanced tradeoff with reduced complexity. When storage is the primary concern, GreedyMinStorage is the best choice, while GreedyLeast-Used is superior when throughput is the sole focus. In future work, we plan to enhance our heuristics by developing a system that selects the most appropriate algorithm based on hardware characteristics and user priorities. Future work will also focus on exploring real-world deployments in storage systems such as Gluster. By advancing the state-of-the-art in data storage algorithms and erasure coding, this work lays

the foundation for resilient and high-performance storage leveraging distributed and heterogeneous storage nodes.

#### Acknowledgments

This material is based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Office of Advanced Scientific Computing Research, under Contract DE-AC02-06CH11357. This publication is partially funded by NSF, under Grant CSSI-2411386 and by MICIU/AEI/10.13039/-501100011033/"FEDER Una manera de hacer Europa" under the project PID2022-138050NB-I00. We gratefully acknowledge the support of Chameleon Cloud for providing the computational resources (project CHI-231082) used in this research.

#### References

- Amazon Web Services. n.d.. Amazon Simple Storage Service (S3). https://aws.amazon.com/s3/ Accessed: 2025-05-01.
- [2] Rachana Ananthakrishnan, Kyle Chard, Ian Foster, and Steven Tuecke. 2015. Globus platform-as-a-service for collaborative science applications. *Concurrency and Computation: Practice and Experience* 27, 2 (2015), 290–305. https://doi.org/10.1002/cpe.3262 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/cpe.3262
- [3] Apache Software Foundation. n.d.. HDFS Erasure Coding (Hadoop 3.0.0 Documentation). https://hadoop.apache.org/docs/r3.0.0/hadoopproject-dist/hadoop-hdfs/HDFSErasureCoding.html. Accessed 2025-05-01.
- [4] Backblaze, Inc. n.d.. Hard Drive Test Data. https://www.backblaze.com/ cloud-storage/resources/hard-drive-test-data. Accessed 2025-05-01.
- [5] Yanpei Chen, Archana Ganapathi, Rean Griffith, and Randy Katz. 2011. The case for evaluating MapReduce performance using workload suites. In IEEE 19th Annual International Symposium on Modelling, Analysis, and Simulation of Computer and Telecommunication Systems. IEEE, Singapore.
- [6] Hilmi Egemen Ciritoglu, John Murphy, and Christina Thorpe. 2019. HaRD: a heterogeneity-aware replica deletion for HDFS. *Journal of big data* (2019).
- [7] Kellie Corona, Katie Osterdahl, Roderic Collins, and Anthony Hoogs. 2021. MEVA: A Large-Scale Multiview, Multimodal Video Dataset for Activity Detection. In *IEEE/CVF Winter Conference on Applications of Computer Vision*.
- [8] Austin Coursey, Gopal Nath, Srikanth Prabhu, and Saptarshi Sengupta. 2021. Remaining Useful Life Estimation of Hard Disk Drives using Bidirectional LSTM Networks. In 2021 IEEE International Conference on Big Data (Big Data). 4832–4841. https://doi.org/10.1109/BigData52589. 2021.9671605
- [9] DAOS Stack. 2024. DAOS Object Module README. https://github. com/daos-stack/daos/blob/v2.6.0/src/object/README.md. Version 2.6.0, accessed 2025-05-01.
- [10] Ishan Dave, Zacchaeus Scheffer, Akash Kumar, Sarah Shiraz, Yogesh Singh Rawat, and Mubarak Shah. 2022. GabriellaV2: Towards Better Generalization in Surveillance Videos for Action Detection. In IEEE/CVF Winter Conference on Applications of Computer Vision Workshops.
- [11] Trinh Viet Doan, Yiannis Psaras, Jörg Ott, and Vaibhav Bajpai. 2022. Toward decentralized cloud storage with IPFS: Opportunities, challenges, and future considerations. *IEEE Internet Computing* (2022).

D-Rex: Heterogeneity-Aware Reliability Framework and Adaptive Algorithms for Distributed Storage

ICS '25, June 08-11, 2025, Salt Lake City, UT, USA

- [12] Matthias Drusch, Umberto Del Bello, Sébastien Carlier, Olivier Colin, Veronica Fernandez, Ferran Gascon, Bianca Hoersch, Claudia Isola, Paolo Laberinti, Philippe Martimort, et al. 2012. Sentinel-2: ESA's optical high-resolution mission for GMES operational services. *Remote Sensing of Environment* 120 (2012), 25–36.
- [13] Jon Elerath and Michael Pecht. 2008. A highly accurate method for assessing reliability of redundant arrays of inexpensive disks (RAID). *IEEE Trans. Comput.* 58, 3 (2008), 289–299.
- [14] Ohad Eytan, Danny Harnik, Effi Ofer, Roy Friedman, and Ronen Kat. 2020. It's Time to Revisit LRU vs. FIFO. In 12th USENIX Workshop on Hot Topics in Storage and File Systems.
- [15] Mikaila J. Gossman, Bogdan Nicolae, and Jon C. Calhoun. 2024. Scalable I/O aggregation for asynchronous multi-level checkpointing. *Future Generation Computer Systems* 160 (2024), 420–432. https://doi.org/10. 1016/j.future.2024.06.003
- [16] Qiang He, Guobiao Zhang, Jiawei Wang, Ruikun Luo, Xiaohai Dai, Yuchong Hu, Feifei Chen, Hai Jin, and Yun Yang. 2024. EdgeHydra: Fault-Tolerant Edge Data Distribution Based on Erasure Coding. *IEEE Transactions on Parallel and Distributed Systems* (2024).
- [17] Michael Hennecke. 2023. Understanding DAOS Storage Performance Scalability. In *HPC Asia 2023 Workshops* (Raffles Blvd, Singapore). Association for Computing Machinery, New York, NY, USA. https: //doi.org/10.1145/3581576.3581577
- [18] Yili Hong. 2013. On computing the distribution function for the Poisson binomial distribution. *Computational Statistics & Data Analysis* (2013). https://doi.org/10.1016/j.csda.2012.10.006
- [19] Hai Jin, Ruikun Luo, Qiang He, Song Wu, Zilai Zeng, and Xiaoyu Xia. 2023. Cost-Effective Data Placement in Edge Storage Systems With Erasure Code. *IEEE Transactions on Services Computing* (2023). https://doi.org/10.1109/TSC.2022.3152849
- [20] Kate Keahey, Jason Anderson, Zhuo Zhen, Pierre Riteau, Paul Ruth, Dan Stanzione, Mert Cevik, Jacob Colleran, Haryadi S Gunawi, Cody Hammock, et al. 2020. Lessons learned from the Chameleon testbed. In USENIX Annual Technical Conference.
- [21] John Kubiatowicz, David Bindel, Yan Chen, Steven Czerwinski, Patrick Eaton, Dennis Geels, Ramakrishna Gummadi, Sean Rhea, Hakim Weatherspoon, Westley Weimer, et al. 2000. OceanStore: An architecture for global-scale persistent storage. ACM SIGOPS Operating Systems Review 34, 5 (2000), 190–201.
- [22] Gil Jae Lee and José AB Fortes. 2019. Improving data-analytics performance via autonomic control of concurrency and resource units. ACM Transactions on Autonomous and Adaptive Systems (2019).
- [23] Ojus Thomas Lee, S. D Madhu Kumar, and Priya Chandran. 2016. Erasure coded storage systems for cloud storage — challenges and opportunities. In *International Conference on Data Science and Engineering*. https://doi.org/10.1109/ICDSE.2016.7823943
- [24] Jun Li and Baochun Li. 2021. Demand-Aware Erasure Coding for Distributed Storage Systems. *IEEE Transactions on Cloud Computing* (2021). https://doi.org/10.1109/TCC.2018.2885306
- [25] Xiaolu Li, Keyun Cheng, Zhirong Shen, and Patrick P. C. Lee. 2022. Fast Proactive Repair in Erasure-Coded Storage: Analysis, Design, and Implementation. *IEEE Transactions on Parallel and Distributed Systems* 33, 12 (2022), 3400–3414. https://doi.org/10.1109/TPDS.2022.3152817
- [26] Zijian Li and Chuqiao Xiao. 2021. ER-Store: A Hybrid Storage Mechanism with Erasure Coding and Replication in Distributed Database Systems. *Scientific Programming* 2021, 1 (2021), 9910942.
- [27] Shu Liu, Xiangxi Mo, Moshik Hershcovitch, Henric Zhang, Audrey Cheng, Guy Girmonsky, Gil Vernik, Michael Factor, Tiemo Bang, Soujanya Ponnapalli, Soujanya Ponnapalli, Natacha Crooks, Joseph E. Gonzalez, Danny Harnik, and Ion Stoica. 2025. SkyStore: Cost-Optimized Object Storage Across Regions and Clouds. arXiv preprint arXiv:2502.20818 (2025).

- [28] Ricardo Marcelín-Jiménez, Jorge Luis Ramírez-Ortíz, Enrique Rodríguez De La Colina, Michael Pascoe-Chalke, and José Luis González-Compeán. 2020. On the Complexity and Performance of the Information Dispersal Algorithm. *IEEE Access* (2020). https://doi.org/10.1109/ ACCESS.2020.3020501
- [29] Bogdan Nicolae, Adam Moody, Elsa Gonsiorowski, Kathryn Mohror, and Franck Cappello. 2019. VeloC: Towards High Performance Adaptive Asynchronous Checkpointing at Large Scale. In *IPDPS'19: The* 2019 IEEE International Parallel and Distributed Processing Symposium. Rio de Janeiro, Brazil, 911–920.
- [30] Bogdan Nicolae, Adam Moody, Greg Kosinovsky, Kathryn Mohror, and Franck Cappello. 2021. VELOC: VEry Low Overhead Checkpointing in the Age of Exascale. In SuperCheck'21: The First International Symposium on Checkpointing for Supercomputing. Virtual Event. https://arxiv.org/pdf/2103.02131.pdf
- [31] Riccardo Pinciroli, Lishan Yang, Jacob Alter, and Evgenia Smirni. 2020. The life and death of SSDs and HDDs: Similarities, differences, and prediction models.
- [32] Mingyuan Qin and Zheng Fu. 2021. A Heterogeneous Hybrid Storage Method Based on Ceph Erasure Code. In 2021 Asia-Pacific Conference on Communications Technology and Computer Science (ACCTCS). https: //doi.org/10.1109/ACCTCS52002.2021.00044
- [33] Red Hat, Inc. n.d.. Creating Dispersed Volumes (Red Hat Gluster Storage 3.1 Administration Guide). https: //docs.redhat.com/fr/documentation/red\_hat\_gluster\_storage/ 3.1/html/administration\_guide/chap-red\_hat\_storage\_volumescreating\_dispersed\_volumes\_1. Accessed 2025-05-01.
- [34] Yanjing Ren, Yuanming Ren, Xiaolu Li, Yuchong Hu, Jingwei Li, and Patrick PC Lee. 2024. {ELECT}: Enabling Erasure Coding Tiering for {LSM-tree-based} Storage. In 22nd USENIX Conference on File and Storage Technologies (FAST 24). 293–310.
- [35] Dante D. Sanchez-Gallegos, J. L. Gonzalez-Compean, Maxime Gonthier, Valerie Hayot-Sasson, J. Gregory Pauloski, Haochen Pan, Kyle Chard, Jesus Carretero, and Ian Foster. 2024. DynoStore. https://dynostore. github.io/dynostore.
- [36] Dante D. Sanchez-Gallegos, J. L. Gonzalez-Compean, Maxime Gonthier, Valerie Hayot-Sasson, J. Gregory Pauloski, Haochen Pan, Kyle Chard, Jesus Carretero, and Ian Foster. 2025. DynoStore: A wide-area distribution system for the management of data over heterogeneous storage. In 25th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGRID). Accepted for publication at CCGRID 2025.
- [37] Zhirong Shen, Yuhui Cai, Keyun Cheng, Patrick P. C. Lee, Xiaolu Li, Yuchong Hu, and Jiwu Shu. 2025. A Survey of the Past, Present, and Future of Erasure Coding for Storage Systems. ACM Trans. Storage (2025). https://doi.org/10.1145/3708994
- [38] Mika J. Straka. [n. d.]. Poisson Binomial Distribution for Python. https: //github.com/tsakim/poibin
- [39] Sacheendra Talluri, Nikolas Herbst, Cristina Abad, Tiziano De Matteis, and Alexandru Iosup. 2024. ExDe: Design space exploration of scheduler architectures and mechanisms for serverless data-processing. *Future Generation Computer Systems* (2024).
- [40] Hakim Weatherspoon and John D. Kubiatowicz. 2002. Erasure Coding Vs. Replication: A Quantitative Comparison. In *Peer-to-Peer Systems*, Peter Druschel, Frans Kaashoek, and Antony Rowstron (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 328–337.
- [41] Stephen B Wicker and Vijay K Bhargava. 1999. Reed-Solomon codes and their applications. John Wiley & Sons.
- [42] Hai Zhou, Dan Feng, and Yuchong Hu. 2022. Bandwidth-aware scheduling repair techniques in erasure-coded clusters: Design and analysis. *IEEE Transactions on Parallel and Distributed Systems* (2022).