IBM

Zurich Research Laboratory

# Relations Between Entity Sizes and Error-Correction Coding Codewords and Data Loss

Ilias Iliadis
ili@zurich.ibm.com
May 26-30, 2024

CTRQ 2024

IARIA

www.zurich.ibm.com

# Short Résumé

- Position
  - IBM Research - Zurich Laboratory since 1988
- Research interests
  - performance evaluation
  - optimization and control of computer communication networks
  - reliability of storage systems
  - storage provisioning for Big Data
  - cloud infrastructures
  - switch architectures
  - stochastic systems
- Affiliations
  - IARIA Fellow
  - senior member of IEEE
  - IFIP Working Group 6.3
- Education
  - Ph.D. in Electrical Engineering from Columbia University, New York
  - M.S. in Electrical Engineering from Columbia University, New York
  - B.S. in Electrical Engineering from the National Technical University of Athens, Greece

# Data Losses in Storage Systems

- Storage systems suffer from data losses due to
  - component failures
    - disk failures
    - node failures
  - media failures
    - unrecoverable and latent media errors

- Reliability enhanced by a large variety of redundancy and recovery schemes
  - RAID systems (**R**edundant **A**rray of **I**ndependent **D**isks)

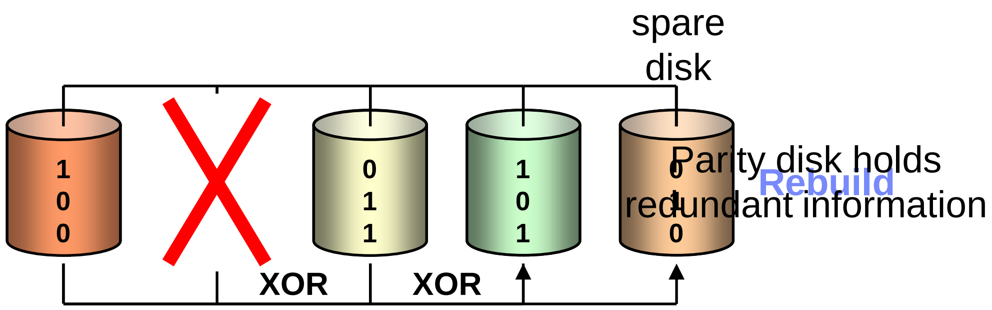


  - RAID-5: Tolerates one disk failure [Patterson *et al*. 1988]

# Data Losses in Storage Systems

- Storage systems suffer from data losses due to
  - component failures
    - disk failures
    - node failures
  - media failures
    - unrecoverable and latent media errors

- Reliability enhanced by a large variety of redundancy and recovery schemes
  - RAID systems

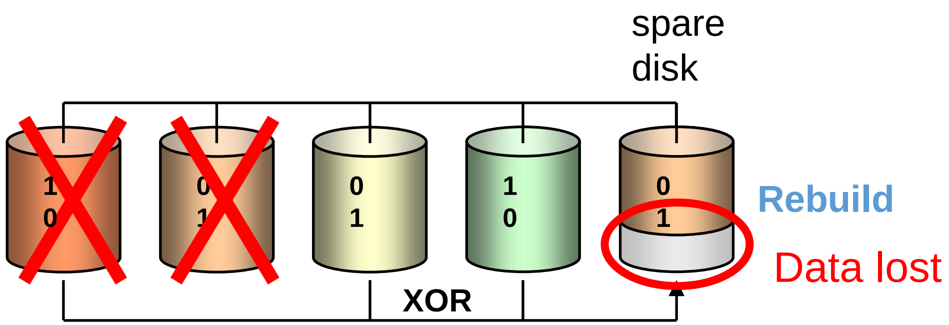


  - RAID-5: Tolerates one disk failure
  - RAID-6: Tolerates two disk failures

# Erasure Coded Schemes

- User data divided into blocks (symbols) of fixed size
  - Complemented with parity symbols
    - codewords

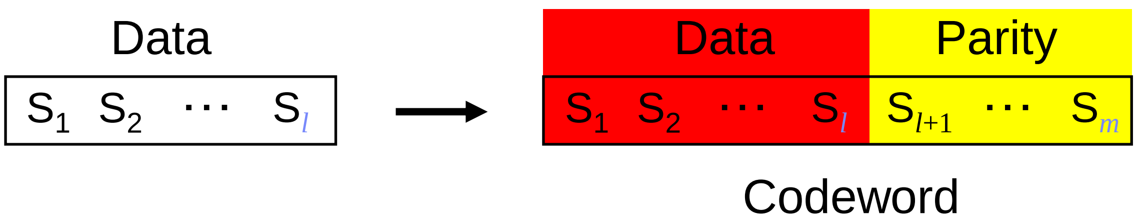


- $(m,l)$ maximum distance separable (MDS) erasure codes
- Any subset of $l$ symbols can be used to reconstruct a codeword
  - Replication : $l = 1$ and $m = r$
  - RAID-5 : $m = l + 1$
  - RAID-6 : $m = l + 2$

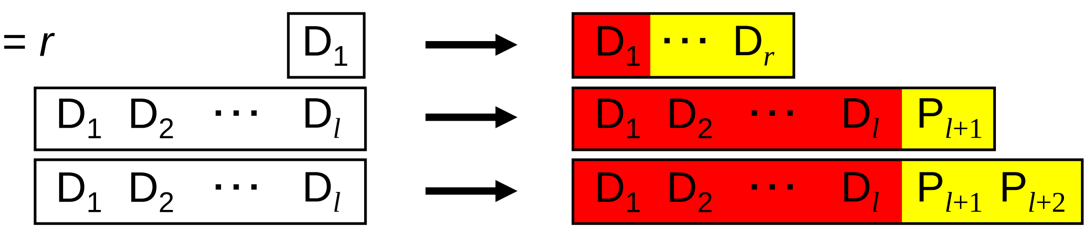


- Storage efficiency : $s_{eff} = l/m$ (Code rate)

- Google : Three-way replication (3,1) → $s_{eff}$ = 33% to Reed-Solomon (9,6) → $s_{eff}$ = 66 %
- Facebook : Three-way replication (3,1) → $s_{eff}$ = 33% to Reed-Solomon (14,10) → $s_{eff}$ = 71 %
- Microsoft Azure : Three-way replication (3,1) → $s_{eff}$ = 33% to LRC (16,12) → $s_{eff}$ = 75 %

# Codeword and Entity Loss

- Erasure coding
  - reduction in storage overhead
  - improvement of reliability achieved

  but
  - repair problem
    - increased network traffic needed to repair data lost
    - Solution: lazy rebuild
      - rebuild process not triggered immediately upon first device failure
      - rebuild process delayed until additional device failures occur
        - ✓ reduces recovery bandwidth
        - ✓ keeps the impact on read performance and data durability low

- **Variable-size entities**
  - each entity spans a number of codewords
  - when a codeword of an entity loses $m - l + 1$ or more symbols, this codeword, and consequently the entity is permanently lost
    - Permanent codeword loss $\Rightarrow$ Permanent entity loss
  - reconstruction of successive codewords leads to the successive reconstruction of entities

# Data Placement of Entities and Formation of Codewords

- Symbol-aligned shards of integer size [ Iliadis, CTRQ 2023 ]

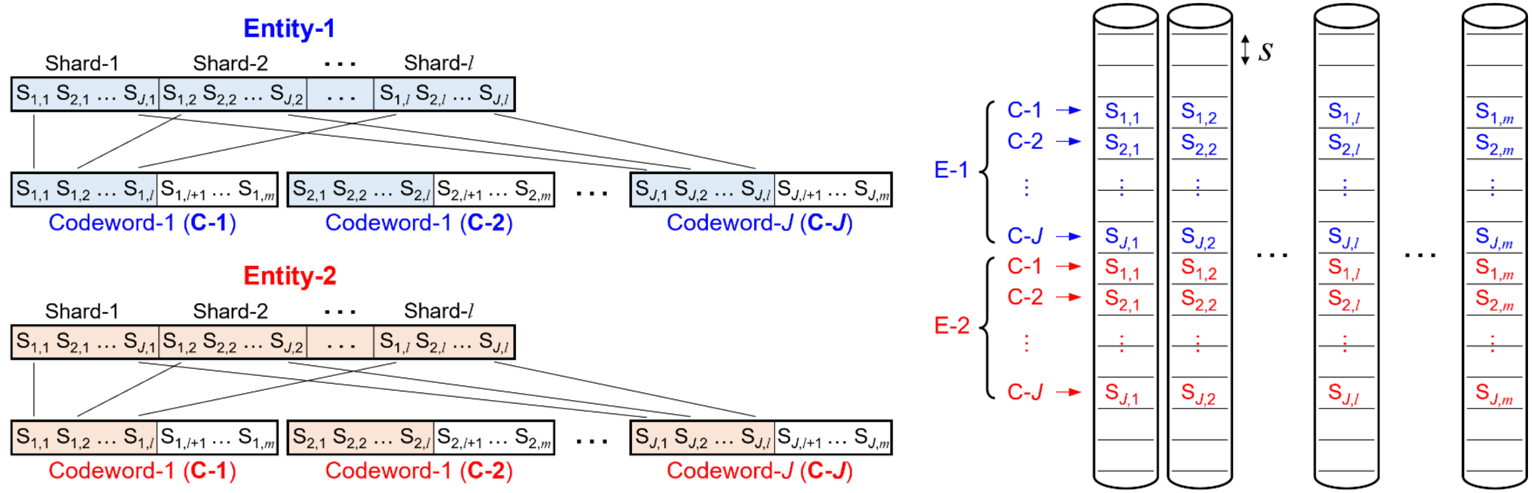


- Non-symbol-aligned shards of arbitrary size

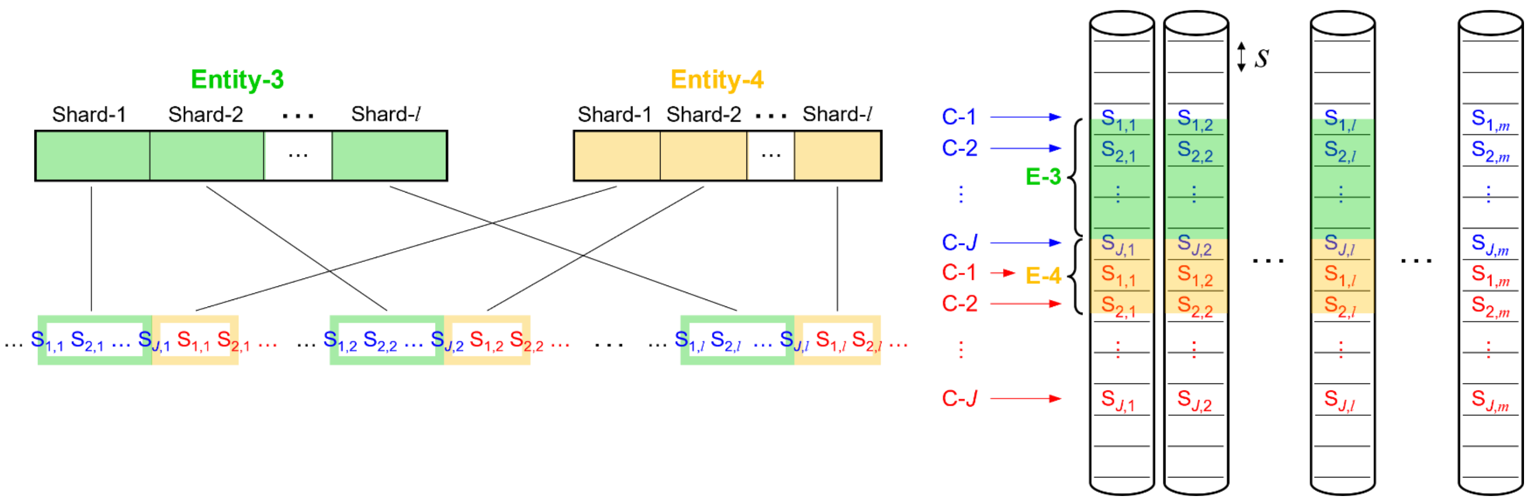

# Reliability Metrics – MTTDL, EAFDL and EAFEL

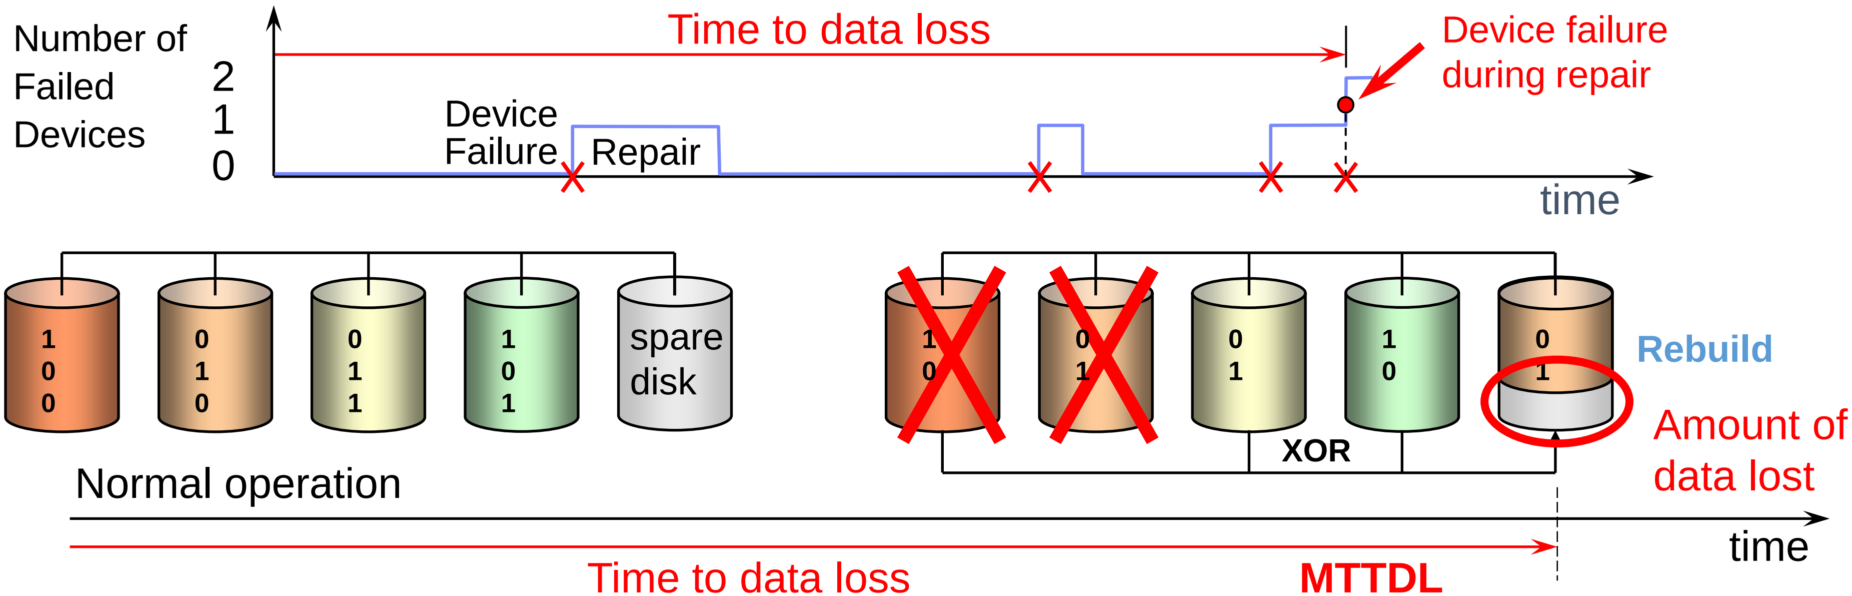


- Data loss events documented in practice by Yahoo!, LinkedIn, Facebook and Amazon
  - Amazon S3 (Simple Storage Service) is designed to provide 99.999999999% durability of objects over a given year
    - average annual expected loss of a fraction of $10^{-11}$ of the data stored in the system
- Assess the implications of system design choices on the
  - frequency of data loss events
    - **Mean Time to Data Loss (MTTDL)**
  - amount of data lost
    - **Expected Annual Fraction of Data Loss (EAFDL)**
      I. Iliadis and V. Venkatesan,
      "Expected Annual Fraction of Data Loss as a Metric for Data Storage Reliability", MASCOTS 2014
    - **Expected Annual Fraction of Entity Loss (EAFEL)**
      I. Iliadis,
      "Expected Annual Fraction of Entity Loss as a Metric for Data Storage Durability", CTRQ 2023

# Reliability of Erasure Coded Systems

- Analytical closed-form expressions for the MTTDL, EAFDL and EAFEL of erasure coded systems in the presence of latent errors when the lazy rebuild scheme is employed
  I. Iliadis, "Effect of Lazy Rebuild on Reliability of Erasure-Coded Storage Systems", CTRQ 2022
  I. Iliadis, "Expected Annual Fraction of Entity Loss as a Metric for Data Storage Durability", CTRQ 2023

- MTTDL does not depend on the placement and size of the entities, but EAFEL does
  - EAFEL metric assesses losses at an entity (file, object, or block) level
  - EAFEL depends on the number of codewords that stored entities span
  - EAFEL reflects the fraction of lost user data only when entities have a fixed size
    - New metric introduced to account for effective user data losses in the case of variable-size entities
      - **Expected Annual Fraction of Effective Data Loss (EAFEDL)**
        - fraction of stored user data that is expected to be lost by the system annually at the entity level

## OBJECTIVE

To derive the distribution of the number of codewords that entities span

To theoretically evaluate the Expected Annual Fraction of Effective Data Loss (EAFEDL)

## RESULTS

- Distribution of the number of codewords that entities span depends on
  - statistics (size and frequency of occurrence) and placement of entities stored
- Evaluation of EAFEL and EAFEDL for variable-size entities

# Symbols Spanned by Shards

- Alternating shard placement of variable-size entities

Shard sequence: {0.3, 2.7, 0.3, 2.7, . . . }

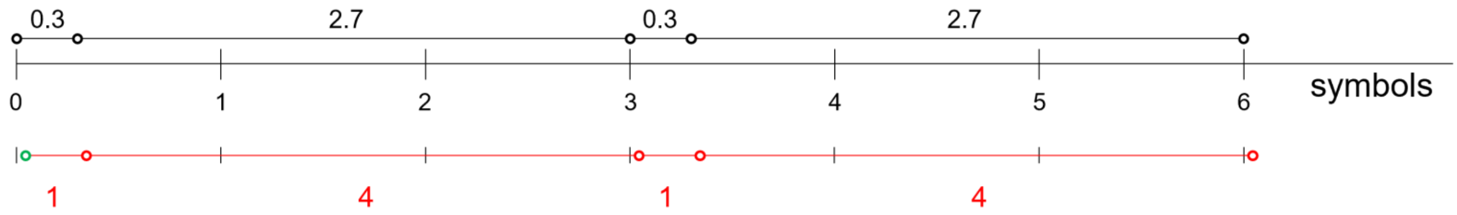


$K$ : number of symbols spanned

$$P(K=i)=p_i=\begin{cases}0.5\,, & \text{for } i=1\\ 0.5\,, & \text{for } i=4\end{cases}$$

Shard sequence: {2.7, 0.3, 2.7, 0.3, . . . }

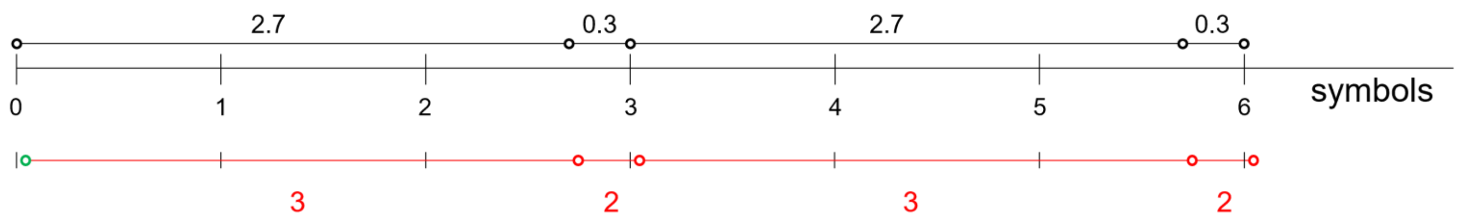


$$P(K=i)=p_i=\begin{cases}0.5\,, & \text{for } i=2\\ 0.5\,, & \text{for } i=3\end{cases}$$

- pdf of $K$ depends on the actual placement

# Symbols Spanned by Randomly Placed Shards

- Notation
  - $l$ : number of user-data symbols per codeword ( $l \geq 1$ )
  - $m$ : total number of symbols per codeword ( $m > l$ )
  - $(m,l)$ : MDS-code structure
  - $e_s$ : entity size, entities of $L$ different sizes: $e_{s,1} < e_{s,2} < \cdots < e_{s,L}$
  - $v_s$ : pdf of entity and shard size
  - $s$ : symbol size
  - $J$ : shard size in symbol-size units $J_j = \frac{e_{s,j}}{l\, s}$ for $j = 1, 2, \ldots, L$
  - $K$ : codewords (symbols) spanned by a shard entity

$$P(K = i) = p_i = \begin{cases} [1 - fr(J_j)]\, v_j\,, & \text{for } i = \lfloor J_j \rfloor + 1 \\ fr(J_j)\, v_j\,, & \text{for } i = \lfloor J_j \rfloor + 2 \\ 0\,, & \text{otherwise}\,, \end{cases} \qquad \text{for } j = 1, 2, \ldots, L$$

where $fr(x) \triangleq x - \lfloor x \rfloor$ denotes the fractional part of $x$

# Non-Markov Analysis for MTTDL, EAFEL, and EAFEDL

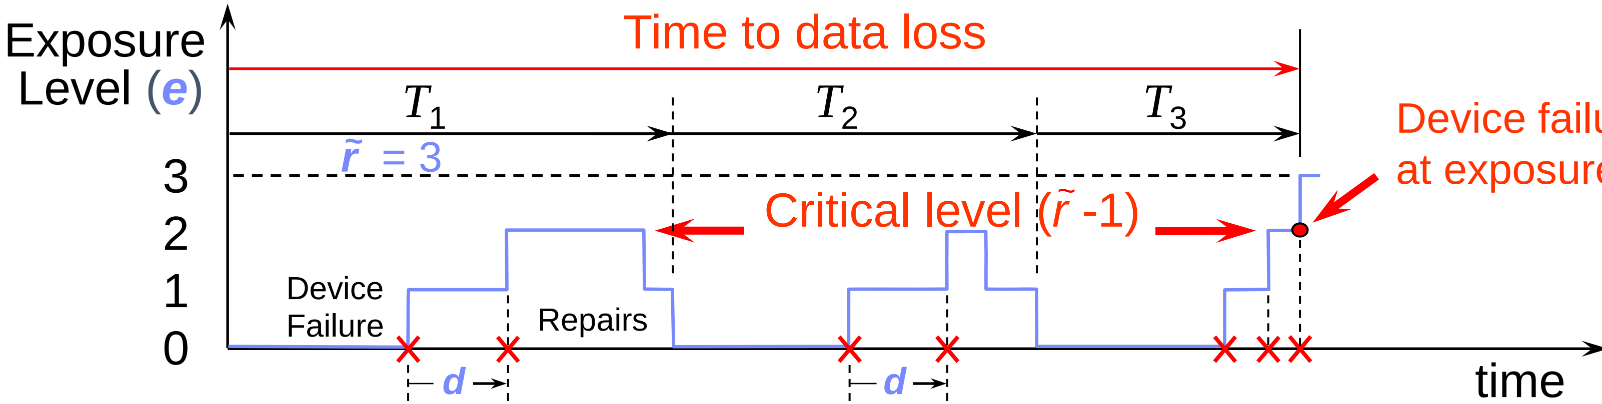


- EAFEL evaluated in parallel with MTTDL
  - $\tilde{r}$ : Minimum number of device failures that may lead to data loss ( $\tilde{r} = m - l + 1$ )
  - $d$ : Lazy rebuild threshold ( $0 \le d < m - l$ )
  - $e$ : Exposure Level: maximum number of symbols that any codeword has lost
  - $T_i$ : Cycles (Fully Operational Periods / Repair Periods)
  - $P_{DL}$ : Probability of data loss during repair period
  - $Y$ : Number of lost entities upon a first-device failure
  - $J$ : Number of codewords per entity
  - $N_E$ : Number of entities stored in a system comprised of $n$ devices
  - $1/\lambda$ : Mean Time to Failure (MTTF) of a device

  - MTTDL $= \sum_i E(T_i) = \dfrac{E(T)}{P_{DL}}$  EAFEL $\approx \dfrac{E(Y)}{E(T)\, NE}$  EAFEDL $\approx \dfrac{m\, E(\check{Q})}{n\, l\, c\, E(T)}$

- System evolution does not depend only on the latest state, but on the entire path
  - underlying models are not semi-Markov

**MTTDL and EAFEL expressions obtained using non-Markov analysis**

# Theoretical Results

- $n$ : number of storage devices
- $k$ : group size (number of devices in a group)
- $c$ : amount of data stored on each device
- $(m,l)$ : MDS erasure code
- $d$ : lazy rebuild threshold
- $b$ : reserved rebuild bandwidth per device
- $B_{max}$ : Maximum network rebuild bandwidth per group of devices
- $1/\lambda$ : mean time to failure of a storage device
- $P_s$ : probability of an unrecoverable sector (symbol) error

$$\text{EAFEL} \approx \frac{E(Y)}{E(T) \cdot N_E} \qquad \text{EAFEDL} \approx \frac{m\, E(\breve{Q})}{n\, l\, c\, E(T)}$$

where

$$E(Y) \approx E(Y_{\text{DF}}) + \sum_{u=d+1}^{\tilde{r}-1} E(Y_{\text{UF}_u})$$

$$E(Y_{\text{UF}_u}) \approx \frac{C}{E(J)} \, \frac{P_u}{u-d} \left( \prod_{j=1}^{u-1} V_j \right) \tilde{q}_u$$

$$E(Y_{\text{DF}}) \approx \frac{C}{E(J)} \, \frac{P_{\text{DF}}}{\tilde{r}-d} \prod_{j=1}^{\tilde{r}-1} V_j$$

$$E(\breve{Q}) \approx E(\breve{Q}_{\text{DF}}) + \sum_{u=d+1}^{\tilde{r}-1} E(\breve{Q}_{\text{UF}_u})$$

$$E(\breve{Q}_{\text{UF}_u}) \approx \frac{C}{E(J)} \, \frac{P_u}{u-d} \left( \prod_{j=1}^{u-1} V_j \right) \breve{q}_u$$

$$E(\breve{Q}_{\text{DF}}) \approx \frac{C}{E(J)} \, \frac{P_{\text{DF}}}{\tilde{r}-d} \left( \prod_{j=1}^{\tilde{r}-1} V_j \right) \breve{q}_{\tilde{r}}$$

$$\tilde{q}_u = \sum_{j=1}^{L} \tilde{q}_{s,u} \left( \frac{e_{s,j}}{l\, s} \right) v_j \qquad N_E \approx \frac{n}{m} \cdot \frac{c}{E(J)\, s}$$

$$\tilde{q}_{s,u}(x) \triangleq 1 - [1 - fr(x)]\, q_u^{f_{\text{cor}}(\lfloor x \rfloor + 1)} - fr(x)\, q_u^{f_{\text{cor}}(\lfloor x \rfloor + 2)}$$

$$q_u = 1 - \sum_{j=\tilde{r}-u}^{m-u} \binom{m-u}{j} P_s^j (1 - P_s)^{m-u-j}$$

$$\breve{q}_u = \sum_{j=1}^{L} e_{s,j}\, \tilde{q}_{s,u} \left( \frac{e_{s,j}}{l\, s} \right) v_j$$

# Numerical Results

- $n$ = 64 : number of storage devices
- $c$ = 20 TB : amount of data stored on each device
- $s$ = 512 B, 5 MB : sector size
- $1/\lambda$ = 876,000 h : MTTF
- $b$ = 100 MB/s : reserved rebuild bandwidth
  - $1/\mu = c/b$ = 55.5 h : MTTR
  - $\lambda/\mu$ = $6\text{x}10^{-5} \ll 1$ : MTTR to MTTF ratio
- $m$ = 16 : number of symbols per codeword
- $P_s$ : $P$ (unrecoverable sector error)

- Numerical results for two system configurations
  - Declustered placement
    - $k = n = 64$
  - Clustered placement
    - $k = 16$
      - System comprises 4 clustered groups

# Effect of Latent Errors on EAFEL and EAFEDL for Declustered Placement

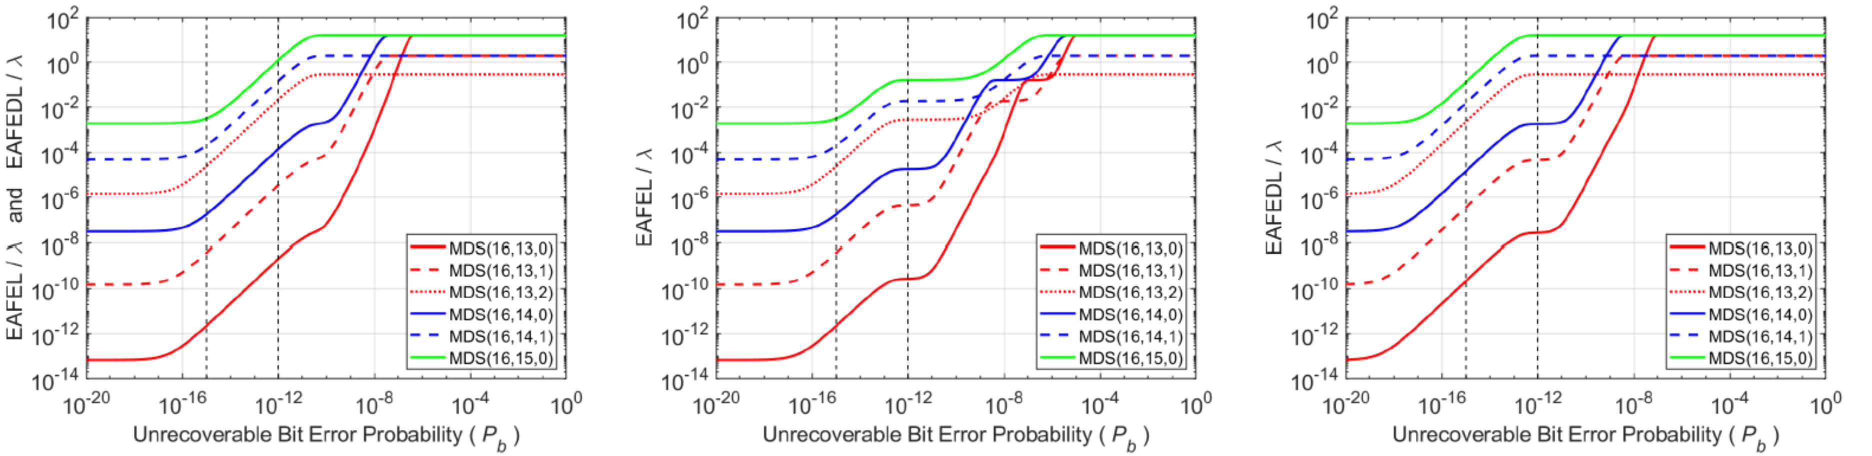


(a) Fixed Entity Size: $e_s = 10$ GB (b) Bimodal Entity Sizes: $e_{s,1} = 1$ MB, $e_{s,2} = 1$ TB (c) Bimodal Entity Sizes: $e_{s,1} = 1$ MB, $e_{s,2} = 1$ TB

- Symbol size of 512 B
- EAFEL and EAFEDL degrade in the interval $[10^{-15}, 10^{-12}]$ of practical interest owing to latent errors
- For fixed size entities, EAFEL and EAFEDL are the same
- Discrete bimodal distribution with average entity size 10 GB
- For large values of $P_b$, EAFEL is reduced whereas EAFEDL is increased
- Increasing the number of parities (reducing $l$ ) improves reliability by orders of magnitude
- Employing lazy rebuild degrades reliability by orders of magnitude
- The declustered placement scheme achieves a significantly lower EAFEL and EAFEDL than the clustered one

# Effect of Latent Errors on EAFEL and EAFEDL for Clustered Placement

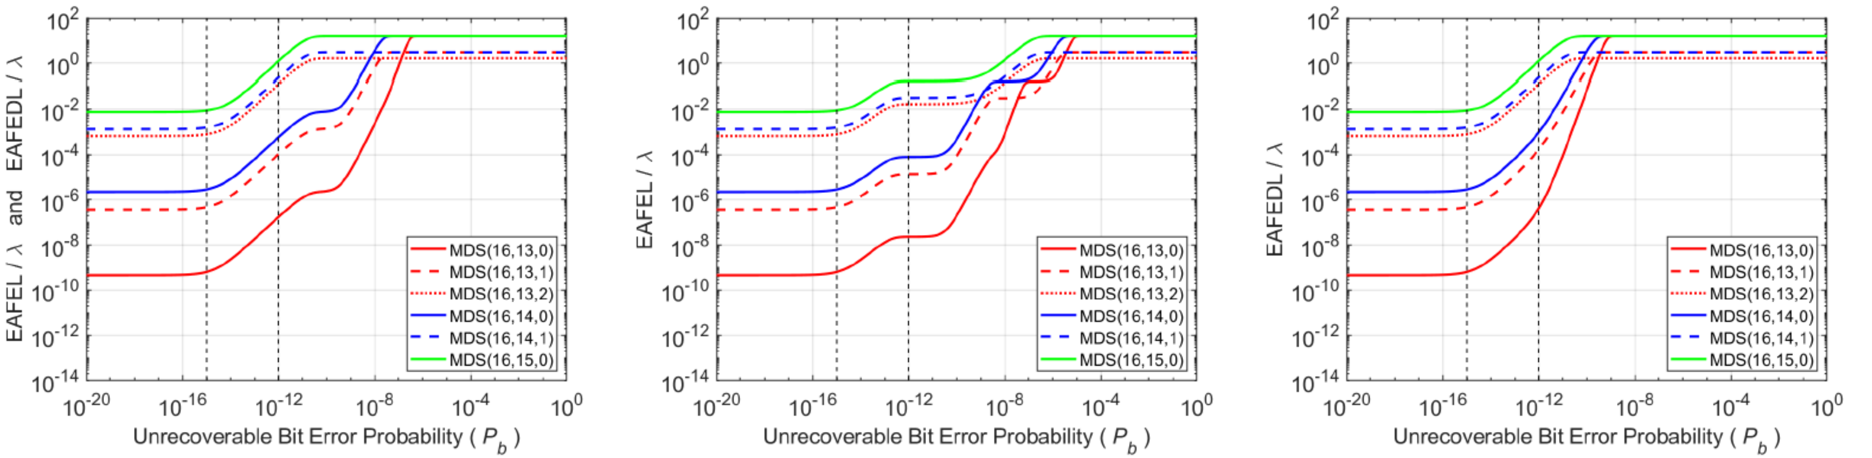


(a) Fixed Entity Size: $e_s = 10$ GB (b) Bimodal Entity Sizes: $e_{s,1} = 1$ MB, $e_{s,2} = 1$ TB (c) Bimodal Entity Sizes: $e_{s,1} = 1$ MB, $e_{s,2} = 1$ TB

- Symbol size of 512 B
- EAFEL and EAFEDL degrade in the interval $[10^{-15}, 10^{-12}]$ of practical interest owing to latent errors
- For fixed size entities, EAFEL and EAFEDL are the same
- Discrete bimodal distribution with average entity size 10 GB
- For large values of $P_b$, EAFEL is reduced whereas EAFEDL is increased
- Increasing the number of parities (reducing $l$ ) improves reliability by orders of magnitude
- Employing lazy rebuild degrades reliability by orders of magnitude
- The clustered placement scheme achieves a significantly higher EAFEL and EAFEDL than the declustered one

 

# Effect of Symbol Size on EAFEL and EAFEDL for Declustered Placement

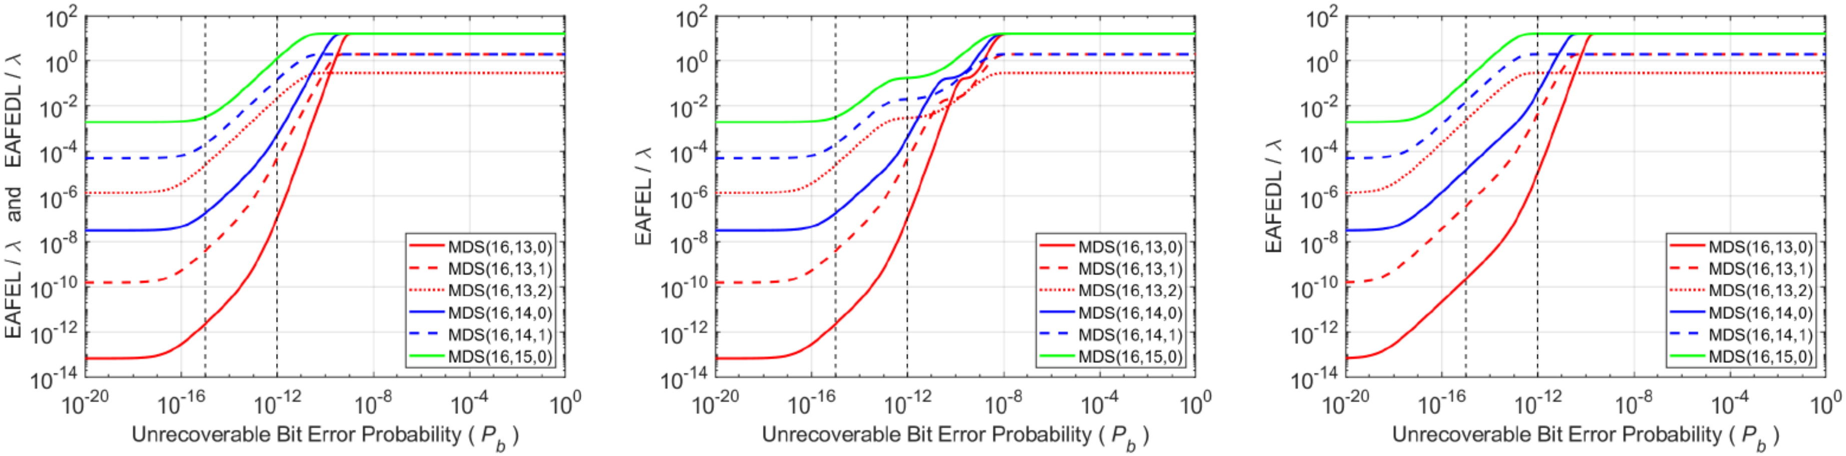


(a) Fixed Entity Size: $e_s = 10$ GB (b) Bimodal Entity Sizes: $e_{s,1} = 1$ MB, $e_{s,2} = 1$ TB (c) Bimodal Entity Sizes: $e_{s,1} = 1$ MB, $e_{s,2} = 1$ TB

- Symbol size of 5 MB
  - EAFEL and EAFEDL degrade compared to the symbol size of 512 B
- EAFEL and EAFEDL degrade in the interval [$10^{-15}$, $10^{-12}$] of practical interest owing to latent errors
- For fixed size entities, EAFEL and EAFEDL are the same
- Discrete bimodal distribution with average entity size 10 GB
- For large values of $P_b$, EAFEL is reduced whereas EAFEDL is increased
- Increasing the number of parities (reducing $l$ ) improves reliability by orders of magnitude
- Employing lazy rebuild degrades reliability by orders of magnitude
- The declustered placement scheme achieves a significantly lower EAFEL and EAFEDL than the clustered one

# Effect of Symbol Size on EAFEL and EAFEDL for Clustered Placement

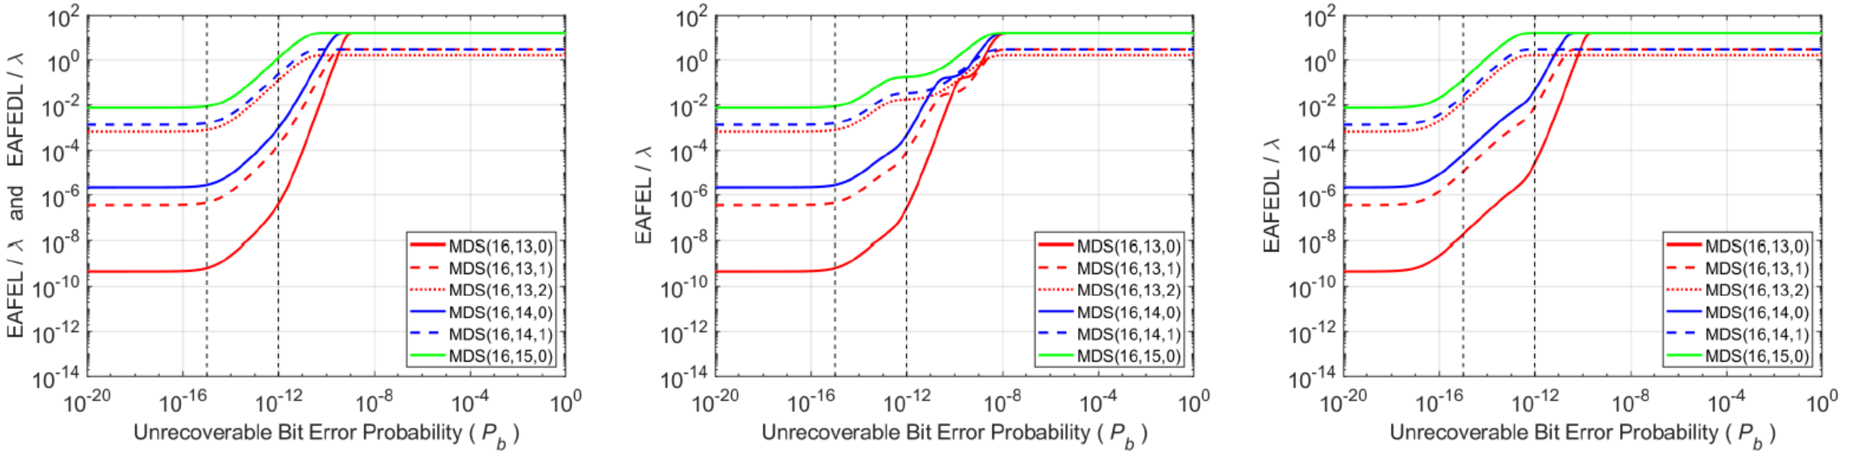


(a) Fixed Entity Size: $e_s = 10$ GB (b) Bimodal Entity Sizes: $e_{s,1} = 1$ MB, $e_{s,2} = 1$ TB (c) Bimodal Entity Sizes: $e_{s,1} = 1$ MB, $e_{s,2} = 1$ TB

- Symbol size of 5 MB
  - EAFEL and EAFEDL degrade compared to the symbol size of 512 B
- EAFEL and EAFEDL degrade in the interval $[10^{-15}, 10^{-12}]$ of practical interest owing to latent errors
- For fixed size entities, EAFEL and EAFEDL are the same
- Discrete bimodal distribution with average entity size 10 GB
- For large values of $P_b$, EAFEL is reduced whereas EAFEDL is increased
- Increasing the number of parities (reducing $l$) improves reliability by orders of magnitude
- Employing lazy rebuild degrades reliability by orders of magnitude
- The clustered placement scheme achieves a significantly higher EAFEL and EAFEDL than the declustered one

 

# CERN File Size Distribution

CERN file size distribution considered in

- I. Iliadis, Y. Kim, S. Sarafijanovic, V. Venkatesan, "Performance Evaluation of a Tape Library System", MASCOTS 2016
- I. Iliadis, L. Jordan, M. Lantz, S. Sarafijanovic, "Performance Evaluation of Automated Tape Library Systems", MASCOTS 2021
- I. Iliadis, L. Jordan, M. Lantz, S. Sarafijanovic, "Performance evaluation of tape library systems", Performance Evaluation 2022

  - mean size: 843 MB
  - second moment: 8.5 $GB^2$
  - standard deviation: 2.8 GB
  - coefficient of variation: 3.4

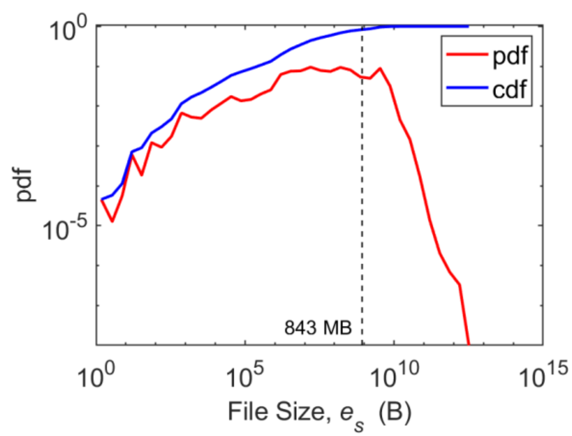


TABLE III. CERN FILE SIZE DISTRIBUTION

| $j$ | Bins | Bin Mean Size $e_{s,j}$ | pdf $v_j$ |
|---|---|---|---|
| 1 | 1 B – 2 B | 2 B | 0.00004559 |
| 2 | 2 B – 5 B | 4 B | 0.00001275 |
| 3 | 5 B – 10 B | 8 B | 0.00005533 |
| 4 | 10 B – 22 B | 16.0 B | 0.00060401 |
| 5 | 22 B – 46 B | 34.0 B | 0.00018569 |
| 6 | 46 B – 100 B | 73.0 B | 0.00121244 |
| 7 | 100 B – 215 B | 157.5 B | 0.00093013 |
| 8 | 215 B – 464 B | 339.5 B | 0.00174431 |
| 9 | 464 B – 1 KB | 732.0 B | 0.00675513 |
| 10 | 1 KB – 2.154 KB | 1.577 KB | 0.00530524 |
| 11 | 2.154 KB – 4.642 KB | 3.398 KB | 0.00496005 |
| 12 | 4.642 KB – 10 KB | 7.321 KB | 0.00800625 |
| 13 | 10 KB – 21.544 KB | 15.772 KB | 0.01174913 |
| 14 | 21.544 KB – 46.416 KB | 33.980 KB | 0.01738480 |
| 15 | 46.416 KB – 100 KB | 73.208 KB | 0.01359001 |
| 16 | 100 KB – 215.443 KB | 157.721 KB | 0.01471745 |
| 17 | 215.443 KB – 464.159 KB | 339.801 KB | 0.02018806 |
| 18 | 464.159 KB – 1 MB | 732.079 KB | 0.02566358 |
| 19 | 1 MB – 2.154 MB | 1.577 MB | 0.06221012 |
| 20 | 2.154 MB – 4.642 MB | 3.398 MB | 0.07519022 |
| 21 | 4.642 MB – 10 MB | 7.321 MB | 0.07654035 |
| 22 | 10 MB – 21.544 MB | 15.772 MB | 0.09501620 |
| 23 | 21.544 MB – 46.416 MB | 33.980 MB | 0.07847651 |
| 24 | 46.416 MB – 100 MB | 73.208 MB | 0.07416942 |
| 25 | 100 MB – 215.443 MB | 157.721 MB | 0.09371673 |
| 26 | 215.443 MB – 464.159 MB | 339.801 MB | 0.08093624 |
| 27 | 464.159 MB – 1 GB | 732.079 MB | 0.05399279 |
| 28 | 1 GB – 2.154 GB | 1.577 GB | 0.04992384 |
| 29 | 2.154 GB – 4.642 GB | 3.398 GB | 0.08871583 |
| 30 | 4.642 GB – 10 GB | 7.321 GB | 0.03182476 |
| 31 | 10 GB – 21.544 GB | 15.772 GB | 0.00452804 |
| 32 | 21.544 GB – 46.416 GB | 33.980 GB | 0.00146156 |
| 33 | 46.416 GB – 100 GB | 73.208 GB | 0.00017060 |
| 34 | 100 GB – 215.443 GB | 157.721 GB | 0.00001375 |
| 35 | 215.443 GB – 464.159 GB | 339.801 GB | 0.00000206 |
| 36 | 464.159 GB – 1 TB | 732.079 GB | 0.00000069 |
| 37 | 1 TB – 2.154 TB | 1.577 TB | 0.00000033 |
| 38 | 2.154 TB – 4.310 TB | 3.230 TB | 0.00000001 |

# Effect of Latent Errors on EAFEL and EAFEDL for Declustered Placement

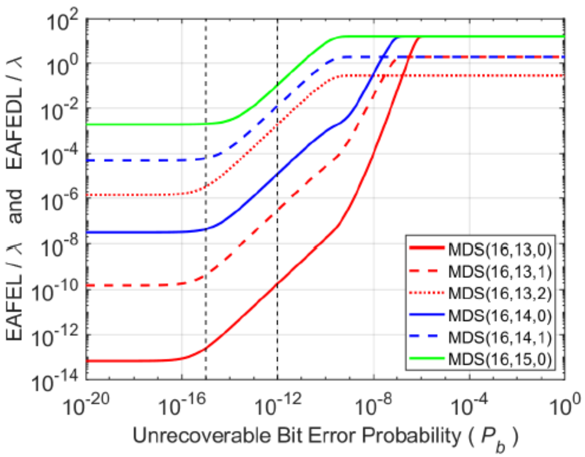


(a) Fixed File Size: $e_s = 843$ MB

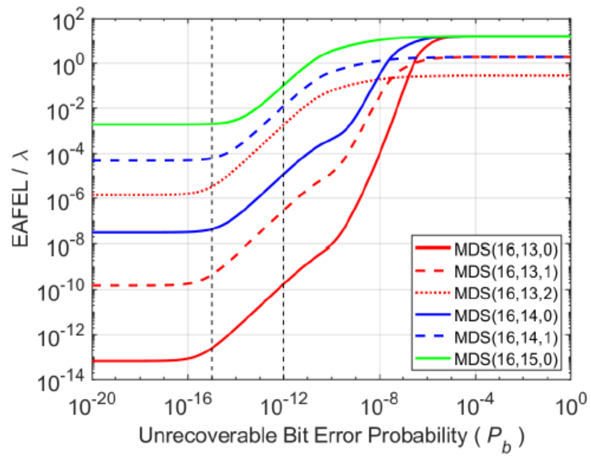


(b) CERN File Sizes; $E(e_s) = 843$ MB

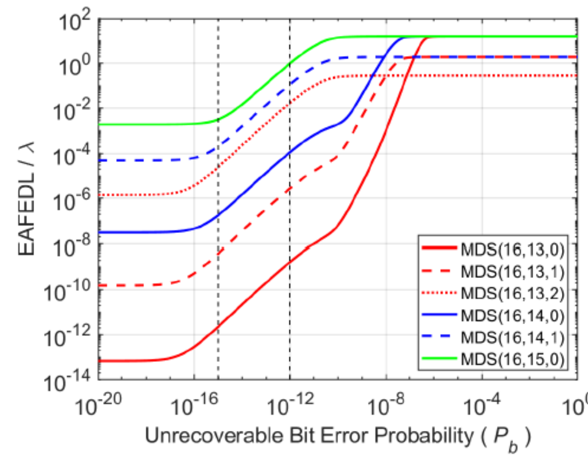


(c) CERN File Sizes; $E(e_s) = 843$ MB

- CERN file size distribution
- Symbol size of 512 B
- EAFEL and EAFEDL degrade in the interval $[10^{-15}, 10^{-12}]$ of practical interest owing to latent errors
- For fixed size entities, EAFEL and EAFEDL are the same
- Discrete bimodal distribution with average entity size 10 GB
- For large values of $P_b$, EAFEL is reduced whereas EAFEDL is increased
- Increasing the number of parities (reducing $l$ ) improves reliability by orders of magnitude
- Employing lazy rebuild degrades reliability by orders of magnitude
- The declustered placement scheme achieves a significantly lower EAFEL and EAFEDL than the clustered one

# Effect of Latent Errors on EAFEL and EAFEDL for Clustered Placement

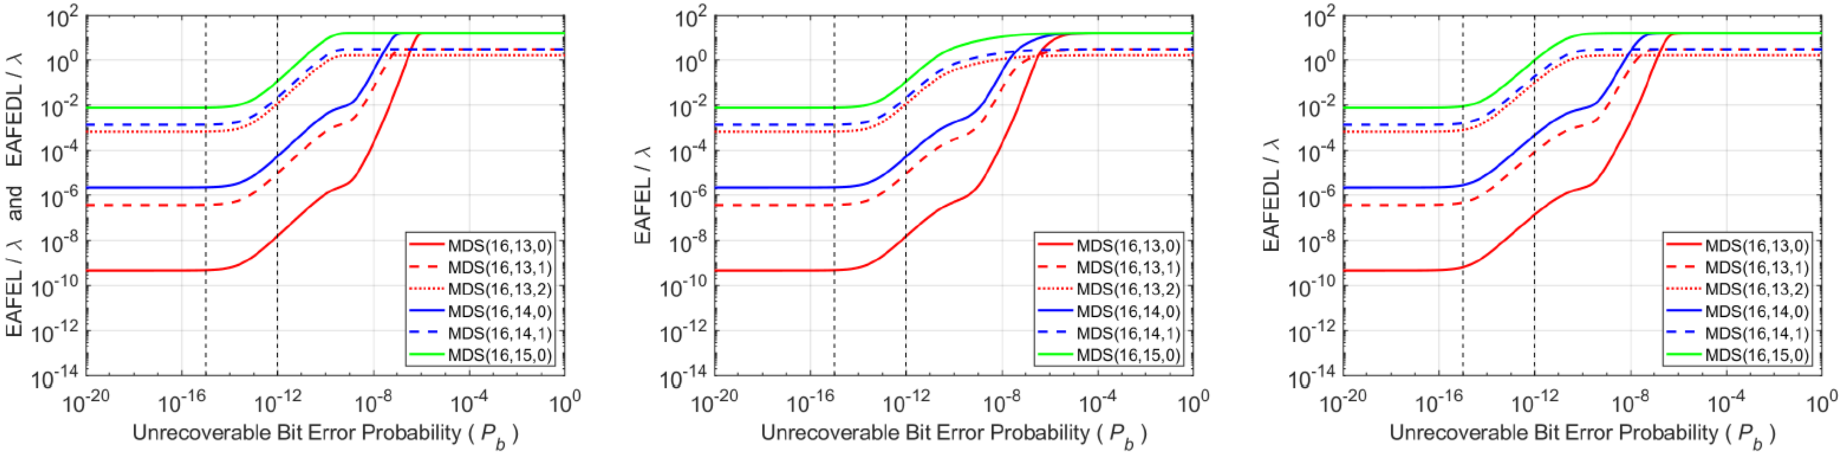


(a) Fixed File Size: $e_s = 843$ MB (b) CERN File Sizes; $E(e_s) = 843$ MB (c) CERN File Sizes; $E(e_s) = 843$ MB

- CERN file size distribution
- Symbol size of 512 B
- EAFEL and EAFEDL degrade in the interval $[10^{-15}, 10^{-12}]$ of practical interest owing to latent errors
- For fixed size entities, EAFEL and EAFEDL are the same
- Discrete bimodal distribution with average entity size 10 GB
- For large values of $P_b$, EAFEL is reduced whereas EAFEDL is increased
- Increasing the number of parities (reducing $l$ ) improves reliability by orders of magnitude
- Employing lazy rebuild degrades reliability by orders of magnitude
- The clustered placement scheme achieves a significantly higher EAFEL and EAFEDL than the declustered one

# Summary

- Introduced the Expected Annual Fraction of Effective Data Loss (EAFEDL) metric, which assesses the durability of distributed and cloud storage systems and reflects losses at an entity (file, object, or block) level
- Considered effect of the lazy rebuild scheme on the reliability of erasure-coded data storage systems
- Assessed the EAFEL and EAFEDL reliability metrics using a non-Markovian analysis
- Derived closed-form expressions for the EAFEL and EAFEDL metrics
- Demonstrated that system reliability is degraded owing to the variability of entity sizes and the employment of the lazy rebuild scheme
- Established that the declustered placement scheme offers superior reliability in terms of both metrics
- Demonstrated that for practical values of unrecoverable sector error probabilities
  - EAFEL and EAFEDL are adversely affected by the presence of latent errors
  - EAFEDL is adversely affected by the entity size variability, but EAFEL improves

# Future Work

- Reliability evaluation of tape storage systems employing erasure-coded schemes

 