

MIT Open Access Articles

The storage vs repair bandwidth trade-off for multiple failures in clustered storage networks

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Abdrashitov, Vitaly et al. "The storage vs repair bandwidth trade-off for multiple failures in clustered storage networks." 2017 IEEE Information Theory Workshop (ITW), November 2017, Kaohsiung, Taiwan, Institute of Electrical and Electronics Engineers (IEEE), February 2018 © 2017 IEEE

Published Version: <http://dx.doi.org/10.1109/itw.2017.8277979>

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Permanent Link: <https://hdl.handle.net/1721.1/121591>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: <http://creativecommons.org/licenses/by-nc-sa/4.0/>



The Storage vs Repair Bandwidth Trade-off for Multiple Failures in Clustered Storage Networks

Vitaly Abdrashitov, N. Prakash and Muriel Médard

Research Laboratory of Electronics, MIT, USA. Email: {vit, prakashn, medard}@mit.edu.

Abstract—We study the trade-off between storage overhead and inter-cluster repair bandwidth in clustered storage systems, while recovering from multiple node failures within a cluster. A cluster is a collection of m nodes, and there are n clusters. For data collection, we download the entire content from any k clusters. For repair of $t \geq 2$ nodes within a cluster, we take help from ℓ local nodes, as well as d helper clusters. We characterize the optimal trade-off under functional repair, and also under exact repair for the minimum storage and minimum inter-cluster bandwidth (MBR) operating points. Our bounds show the following interesting facts: 1) When $t|(m-\ell)$ the trade-off is the same as that under $t = 1$, and thus there is no advantage in jointly repairing multiple nodes, 2) When $t \nmid (m-\ell)$, the optimal file-size at the MBR point under exact repair can be strictly less than that under functional repair. 3) Unlike the case of $t = 1$, increasing the number of local helper nodes does not necessarily increase the system capacity under functional repair.

I. INTRODUCTION

We study the storage-overhead vs repair-bandwidth trade-off for multiple node failures under the setting of clustered storage networks. Our model is motivated by applications to cloud storage settings, where user data is spread across distinct data-centers, even possibly belonging to different service providers (as in a cloud-of-cloud setting). Practical implementation studies that show the benefits of Reed-Solomon codes for data storage in cloud-of-cloud settings appear in [1]–[3]. In our model, a cluster represents a data center. In such networks, it is common to differentiate between intra- and inter-cluster bandwidth costs; typically, intra-cluster bandwidth cost is much less than inter-cluster bandwidth cost. To keep the model simple, we ignore any hierarchical topology that may be present within a data-center (cluster), and simply assume equal cost connectivity between any two nodes inside a cluster. We also assume direct connectivity between any two clusters in the network.

In our model, a cluster is a collection of m physical nodes (see Fig. 1), each of size α symbols from the finite field \mathbb{F}_q , for some q . There are n clusters in total in the system. A file of size B symbols is encoded into $nm\alpha$ symbols and stored across the nm storage nodes. We follow a clustering approach for both data collection and repairs. For data collection, we demand that the entire content of an arbitrary set of k clusters is sufficient to decode the original file. Thus, during data collection, we assume a cluster to be either completely




Fig. 1. System model for clustered data storage, illustrating data collection and repair from multiple failures. The model is suitable for storing coded data across multiple data centers (clusters) as in cloud-of-cloud applications.

Special Case of Our Model	Prior Work
No clustering, single-node repair: $m = 1, t = 1, \ell = 0$	Classical RC, [5]
With clustering, single-node repair: $t = 1$	Generalized RC, [6], Studies impact of ℓ on 1) storage vs Inter-cluster BW trade-off 2) Intra-cluster BW
With clustering, multiple-node repairs: $t \geq 1$	Two Layer coding scheme [7] Study limited to $\ell = m - t$


TABLE I

SPECIAL CASES OF SYSTEM MODEL APPEARING IN LITERATURE.

available (if we connect to it) or completely unavailable (if we do not connect to it). Such an assumption is realistic in a multi-data-center cloud setting [4].

Nodes within a cluster represent failure domains; in this work, we deal with the problem of recovery from t node-failures that occur in one of the n clusters. While single-node failure is the most common failure event, correlated failures of nodes within a data center is an important issue reported in practice [8] and this motivates our failure model. The t *newcomer* nodes are added to the same cluster as replacement to those failed. For restoring the content of the t new nodes, we download local as well as external content. External help is taken from any set of d other clusters, each of which contributes β inter-cluster symbols. The β symbols from any cluster is possibly a function of all the $m\alpha$ symbols in the cluster. For completeness of the model, we assume the presence of a compute unit in the cluster that can combine these $m\alpha$ symbols to generate the β helper symbols. We also download entire content from any set of $\ell \leq m - t$ surviving nodes in the failure cluster. Once again, we assume

This work is in part supported by the Air Force Office of Scientific Research (AFOSR) under award No FA9550-14-1-043, and in part supported by the National Science Foundation (NSF) under Grant No. CCF-1527270.



(a) Trade-off for an $(n = 5, k = 4, d = 4)$ $(m = 3, \ell = 0, t = 2)$ system.

(b) RLNC simulation for an $(n = 3, k = 2, d = 2)$ $(\alpha = 2, \beta = 2)$ $(m = 3, \ell = 0, t = 2)$ system.

(c) Impact of number of local helper nodes, ℓ , on file-size for an $(n = 7, k = 4, d = 5, m = 17, t = 5)$ clustered storage system at MBR point $(\alpha = 1, \beta = 1)$. Local help does not provide any advantage unless $\ell > 2$.

Fig. 2. Illustration of the implications of the exact and functional repair file-size bounds.

the presence of a compute unit in the failure cluster that combines all the local and external helper data, and generates the content of the replacement nodes. We assume that the encoding function does not introduce any local dependence among the nodes of a cluster; for e.g., the model excludes the possibility of a local parity node within a cluster. An analysis of the impact of such local parity nodes is left for future work. We also restrict ourselves to the case $d \geq k$, even though analysis for the case $0 \leq d \leq k - 1$ is perfectly feasible.

A code satisfying the above model requirements for repair and data collection shall be called multi-node repair generalized regenerating code (MRGRC) \mathcal{C} with parameters $\{(n, k, d), (\alpha, \beta), (m, \ell, t)\}$. In this paper, we study the trade-off between storage-overhead (S.O.) $\frac{mn\alpha}{B}$ and inter-cluster (IC) repair-bandwidth-overhead $\frac{d\beta}{t\alpha}$ for the above setting, under both functional and exact repair. Under exact repair, the content of any of the t new nodes is exactly the same as what was stored before failure, while in functional repair, the restored content allows data-collection and further repairs.

Special cases of the model have been studied in the past (see Table I). The setting of regenerating codes (RC) introduced in [5] corresponds to the case with $t = 1, m = 1, \ell = 0$ - we refer to these as classical RCs. The case of single node failure ($t = 1$) in clustered systems was previously studied in [6], where the authors first identify the storage-vs-inter-cluster-BW trade-off (ignoring intra-cluster BW), and then find bounds on the minimum intra-cluster BW that is needed to achieve this trade-off. The authors, show the surprising fact that while increasing the number of local helper nodes ℓ improves the storage-vs-inter-cluster-repair-BW trade-off; increasing ℓ not only increases intra-cluster BW in the host-cluster (this is obvious since one downloads $\ell\alpha$ symbols), but also increases the intra-cluster BW in the d remote helper clusters. In other words, in situations when intra-cluster BW cannot be entirely ignored, the choice of the number of local nodes becomes an important one.

Motivated by the above result of [6], for the case of multiple failures that we consider here, even though we do not explore bounds on intra-cluster BW in this paper, we still parametrize the number of local helper nodes ℓ in the range $0 \leq \ell \leq m - t$, so that our results remain relevant for a future study on intra-

cluster BW for $t \geq 1$. The case of $\ell = m - t, t \geq 1$ has been previously studied in [7]. However, as we show in this paper, even when restricted to storage-vs-inter-cluster-BW trade-off (ignoring intra-cluster BW), the case $0 \leq \ell < m - t, t > 1$ offers several surprising results which cannot be inferred from an analysis of the case $\ell = m - t, t > 1$. Following is a summary of our results in this paper:

A. Our Results

(a) File-Size bound under functional repair: Let $(m - \ell) = at + b, a \geq 1, 0 \leq b \leq t - 1$. Then, the file-size under functional repair of an MRGRC is upper bounded by $B \leq B_F^*$, where

$$B_F^* = \ell k \alpha + a \sum_{i=0}^{k-1} \min(t\alpha, (d-i)\beta) + \sum_{i=0}^{k-1} \min(b\alpha, (d-i)\beta). \quad (1)$$

The bound is shown by considering the information-flow graph (IFG) under functional-repair, and calculating the minimum cut. The bound is indeed tight, if there is a known upper bound on the number of repairs in the system - the achievability follows from results in network coding [9].

(b) File-Size bound under exact repair: For exact repair, we prove a tighter bound, given by

$$B \leq B_E^* = \ell k \alpha + (m - \ell) \sum_{i=0}^{k-1} \min\left(\alpha, \frac{(d-i)\beta}{t}\right). \quad (2)$$

We note that $B_E^* \leq B_F^*$. The bound is optimal at the minimum storage-overhead (MSR) and the minimum inter-cluster repair-bandwidth-overhead (MBR) points characterized by $B = mk\alpha$ and $t\alpha = d\beta$, respectively. We show how optimal constructions for the case $t > 1$ can be directly obtained from optimal constructions for the case $t = 1$ [6].

Implications of the Bounds: *Case a) $t|(m - \ell)$:* In this case, the bounds in (1) and (2) coincide. Further, (1) gives the same S.O. vs IC-repair-bandwidth-overhead trade-off for any value of $t \geq 1$. i.e., under functional repair, there is no advantage to jointly repairing multiple nodes (instead of repairing one). For exact repair, at the MSR and MBR points, there is no benefit to jointly repairing multiple nodes for any $t > 1$, irrespective of if $t|(m - \ell)$ or not.

Case b) $t \nmid (m - \ell)$: In this case, it is possible that $B_F^* > B_E^*$. Specifically, at the MBR point with $t\alpha = d\beta$, we have $B_F^* > B_E^*$, whenever $k > 1$. This also means that the S.O. vs IC-repair-bandwidth-overhead trade-off under functional repair for the case $t > 1$ (with $k > 1$) is strictly better than that for the case $t = 1$. A comparison of trade-offs between exact and functional repair for the case of $\{(n = 5, k = 4, d = 4)(m = 3, \ell = 0, t = 2)\}$ is shown in Fig. 2a. In Fig. 2b, we present a simulation result that shows probability of successful decoding while using random linear network codes (RLNCs) [10] with sufficiently large field size in an $\{(n = 3, k = 2, d = 2)(m = 3, \ell = 0, t = 2)\}$ storage system operating at the MBR point with $\beta = 2$. In this case, optimal file-sizes under exact and functional repair are $B_E^* = 9$ and $B_F^* = 10$. RLNCs enable functional repair, and our simulation result indeed confirms the achievability of file-size $B_F^* = 10$.

Another implication of the bounds relates to the usefulness of the number of local helper nodes ℓ used in the repair process. Under functional repair, for the case of $t = 1$ [6], if we fix $n, k, d, m, t, \alpha, \beta$, the optimal file-size increases strictly monotonically with ℓ , whenever $\alpha > (d - k + 1)\beta$ (i.e., if we exclude the MSR point). However, strict monotonicity is not necessarily true when $t > 1$. Specifically, at the MBR point, it can be shown that whenever $(m \bmod t) \leq \lfloor (d - k + 1)t/d \rfloor$, for any ℓ in the range $0 \leq \ell \leq (m \bmod t)$, the capacity is as good as with no local help at all (see Fig. 2c).

B. Other Related Work

The problem of multiple-node repair for classical RCs have been studied under the frameworks of cooperative repair [11], [12] and centralized repair [13], [14]. In cooperative repair, each of the t replacement nodes first individually contacts respective sets of d helper nodes, and then communicates among themselves before restoring the new content. In centralized repair, a centralized compute node downloads data from some subset of d nodes, and generates the data for all t replacement nodes. Our repair model can be considered as a centralized repair model for clustered storage systems.

Regenerating code variations for data-center-like topologies consisting of racks and nodes are considered in [15]–[19]. All these works focus on single-node repair, whereas we focus on multiple-node repairs. Further, the models in [15], [16] and [17] use clustering approach only for repair (by distinguishing inter and intra rack repair costs), and not for data-collection. File retrievability is demanded from any set of k nodes in the whole system, irrespective of which clusters they belong to. The difference in data collection model is the main difference between our model and the models in [15], [16] and [17].

We next describe how exact repair codes for $t > 1$ can be directly obtained from exact repair codes for $t = 1$. In Sections III and IV we discuss the exact-repair and functional-repair bounds, respectively. For functional repair, our IFG model is substantially different, and more elaborate than the one used in [7] for the case of $\ell = m - t$. The complexity of our model comes from the need to handle the case $\ell < m - t$.

II. EXACT REPAIR CODES

Optimal constructions of exact repair MRGRCs for any $t > 1$ can be directly obtained from constructions for the case $t = 1$, whenever $t|\beta$. In order to construct an exact repair MRGRC \mathcal{C} with parameters $(n, k, d)(\alpha, \beta)(m, \ell, t), t|\beta$, we start with an exact repair code [6] \mathcal{C}' with parameters $\{(n, k, d)(\alpha, \beta' = \beta/t)(m, \ell, t' = 1)\}$. The code \mathcal{C}' was shown to exist at the MSR and MBR points; in fact it was shown in [6] that an optimal $(n, k, d)(\alpha, \beta')(m, \ell, t = 1)$ \mathcal{C}' can be constructed whenever a classical exact repair $(n, k, d)(\alpha, \beta')$ RC exists, with file-size $\sum_{i=0}^{k-1} \min(\alpha, (d - i)\beta')$.

The code \mathcal{C}' can be viewed as the code \mathcal{C} as it is, if we assume that repair of any group of t nodes in \mathcal{C} happens one node at a time via the repair procedure in \mathcal{C}' . Also, we use the same set of local and external helpers for the repair of any of the t failed nodes. Inter-cluster bandwidth, for the repair of the entire group, per external helper amounts to $\beta = t\beta'$. The file-size B that we obtain is given by

$$\begin{aligned} B = B' &= \ell k \alpha + (m - \ell) \sum_{i=0}^{k-1} \min(\alpha, (d - i)\beta') \\ &= \ell k \alpha + (m - \ell) \sum_{i=0}^{k-1} \min\left(\alpha, \frac{(d - i)\beta}{t}\right). \end{aligned} \quad (3)$$

III. FILE SIZE BOUND, EXACT REPAIR

In this section, we present the proof of the file-size upper bound in (2) for exact repair codes. We assume the code to be deterministic; by this we mean that the helper data is uniquely determined given the indices of the t failed nodes, local helper nodes and helper clusters. We begin with useful notation. Let \mathcal{F} denote the random variable corresponding to the data file that gets stored. We assume \mathcal{F} to be uniformly distributed over \mathbb{F}_q^B . Let $Y_{i,j} \in \mathbb{F}_q^\alpha, 1 \leq i \leq n, 1 \leq j \leq m$ denote the content stored in node j of cluster i . For $j \leq j'$, we write $Y_{i,[j,j']}$ to denote $\{Y_{i,j}, Y_{i,j+1}, \dots, Y_{i,j'}\}$. We also write \mathbf{Y}_i to denote $Y_{i,[1,m]}$. Further, for $i \leq i'$, $\mathbf{Y}_{[i,i']}$ will denote $\{\mathbf{Y}_i, \dots, \mathbf{Y}_{i'}\}$. The property of data collection demands that

$$H(\mathcal{F}|\{\mathbf{Y}_i, i \in S\}) = 0 \quad \forall S \subset [n], |S| = k, \quad (4)$$

where $H(\cdot)$ denotes the entropy function computed with respect to $\log q$. Next, consider the repair of t nodes indexed by \mathcal{R}_i in cluster i . Let $\mathcal{H} \subset [n] \setminus \{i\}, |\mathcal{H}| = d$, and $\mathcal{L} \subset [m] \setminus \mathcal{R}_i, |\mathcal{L}| = \ell$ respectively denote the indices of helper clusters and local nodes that aid in the repair process. Let $Z_{i',\mathcal{R}_i}^{\mathcal{H},\mathcal{L}}$ denote helper data passed by cluster i' . The property of exact repair is jointly characterized by the following set of inequalities: $H(Z_{i',\mathcal{R}_i}^{\mathcal{H},\mathcal{L}}|\mathbf{Y}_{i'}) = 0$, $H(Z_{i',\mathcal{R}_i}^{\mathcal{H},\mathcal{L}}) \leq \beta$, and

$$\begin{aligned} H\left(\{Y_{i,j}, j \in \mathcal{R}_i\}|\{Z_{i',\mathcal{R}_i}^{\mathcal{H},\mathcal{L}}, Y_{i,j'}, i' \in \mathcal{H}, j' \in \mathcal{L}\}\right) &= 0, \\ \forall \mathcal{H} \subset [n] \setminus \{i\}, |\mathcal{H}| = d, \forall \mathcal{L} \subset [m] \setminus \mathcal{R}_i, |\mathcal{L}| = \ell. \end{aligned} \quad (5)$$

Our proof technique of the file-size bound presented here, though has some similarity with the information theoretic techniques in works like [14], [20], it differs in an important way. The proofs in these other works rely on the chain rule of entropy, and so does our proof; however in here we demand

that the chain be expanded in a specific order. The following lemma is used to determine this order. The lemma is valid only when $b > 0$, where $(m - \ell) = at + b$, $a \geq 1$, $0 \leq b \leq t - 1$. When $b = 0$, the bound-proof does not need this lemma.

Lemma III.1. *Let $(m - \ell) = at + b$, $a \geq 1$, $1 \leq b \leq t - 1$. Consider any $S_i \subset [n]$, $|S_i| = i$, $1 \leq i \leq k - 1$, and let $\mathcal{Y}(S_i) = \{\mathbf{Y}_i, i \in S_i\}$. Then, for any $i' \in [n] \setminus S_i$, there exists a permutation σ_{i', S_i} of $\{\ell + 1, \ell + 2, \dots, m\}$ such that*

$$H\left(Y_{i', \sigma_{i', S_i}(j')} | \mathcal{Y}(S_i), \tilde{\mathcal{Y}}(i', S_i, j')\right) \leq \min\left(\alpha, \frac{(d - i)\beta}{t}\right), \quad (6)$$

for all $j' \in \{m - b + 1, m - b + 2, \dots, m\}$, where

$$\tilde{\mathcal{Y}}(i', S_i, j') = Y_{i', [1, \ell]} \cup \{Y_{i', \sigma_{i', S_i}(j)}, j \in [\ell + 1, j' - 1]\}. \quad (7)$$

Proof. In here, we only present the candidate for the permutation σ_{i', S_i} . The proof that this satisfies the lemma can be found in Appendix A. Consider the content, of the cluster i' , given by $\{Y_{i', 1}, Y_{i', 2}, \dots, Y_{i', m}\}$. Define the quantities $(j_m, \mathcal{V}_m), (j_{m-1}, \mathcal{V}_{m-1}), \dots, (j_{m-b+1}, \mathcal{V}_{m-b+1})$ in this respective order as below:

Step 1. Let $\mathcal{U} = \{Y_{i', \ell+1}, Y_{i', \ell+2}, \dots, Y_{i', m}\}$, and $x = 0$
 Step 2. Define $(j_{m-x}, \mathcal{V}_{m-x})$ as

$$(j_{m-x}, \mathcal{V}_{m-x}) = \arg \min_{\substack{(j, \mathcal{V}) : \\ Y_{i', j} \in \mathcal{U} \\ \mathcal{V} \subset \mathcal{U} \setminus \{Y_{i', j}\}, |\mathcal{V}| = t-1}} \Theta, \quad (8)$$

where $\Theta = H(Y_{i', j} | \mathcal{V}, \mathcal{Y}(S_i), Y_{i', [1, \ell]})$.

Step 3. If $x < b - 1$, update \mathcal{U} as $\mathcal{U} = \mathcal{U} \setminus \{Y_{i', j_{m-x}}\}$. Increment x by 1 and return to Step 2.

Additionally, let us also define $\{j_{\ell+1}, j_{\ell+2}, \dots, j_{m-b}\} \triangleq \{\ell + 1, \dots, m\} \setminus \{j_m, j_{m-1}, \dots, j_{m-b+1}\}$. In the preceding definition, we only need equality as sets. We do not care about any particular ordering of the elements in $\{\ell + 1, \dots, m\} \setminus \{j_m, j_{m-1}, \dots, j_{m-b+1}\}$, while associating these with $\{j_{\ell+1}, j_{\ell+2}, \dots, j_{m-b}\}$. The candidate for the permutation σ_{i', S_i} on the set $\{\ell + 1, \dots, m\}$ is now defined as follows:

$$\sigma_{i', S_i}(p) = j_p, \quad \ell + 1 \leq p \leq m. \quad (8)$$

□

Proof of Exact Repair Upper Bound: We have

$$\begin{aligned} B &= H(\mathcal{F}) \leq H(\mathbf{Y}_{[1, k]}) = \sum_{i'=1}^k H(\mathbf{Y}_{i'} | \mathbf{Y}_{[1, i'-1]}) \\ &= \sum_{i'=1}^k (H(Y_{i', [1, \ell]} | \mathbf{Y}_{[1, i'-1]}) + H(Y_{i', [\ell+1, m]} | Y_{i', [1, \ell]}, \mathbf{Y}_{[1, i'-1]})) \\ &\leq \ell k \alpha + \sum_{i'=1}^k H(Y_{i', [\ell+1, m]} | Y_{i', [1, \ell]}, \mathbf{Y}_{[1, i'-1]}). \end{aligned} \quad (9)$$

Now, if we let $\sigma = \sigma_{i', [1, i'-1]}$ to be the permutation as obtained from Lemma III.1, then we expand the term




Fig. 3. An illustration of the information flow graph used in cut-set based upper bound for the file-size under functional repair. We assume $(n = 3, k = 2, d = 2)(m = 3, \ell = 0, t = 2)$. Only a subset of nodes are named so as to avoid clutter. Two batches, each of $t = 2$ nodes, fail and get repaired first in cluster 1 and then in cluster 3. We also indicate a possible choice of the $S - T$ cut that results in the desired upper bound. We fail nodes in cluster 3 instead of cluster 2 only to make the figure compact.

$H(Y_{i', [\ell+1, m]} | Y_{i', [1, \ell]}, \mathbf{Y}_{[1, i'-1]})$ in (9) using the order determined by the permutation σ , as follows:

$$\begin{aligned} &H(Y_{i', [\ell+1, m]} | Y_{i', [1, \ell]}, \mathbf{Y}_{[1, i'-1]}) \\ &= H(\{Y_{i', \sigma(j')}, j' \in [\ell + 1, m]\} | Y_{i', [1, \ell]}, \mathbf{Y}_{[1, i'-1]}) \\ &\leq \sum_{u=0}^{a-1} H(\{Y_{i', \sigma(\ell+ut+v)}, v \in [1, t]\} | Y_{i', [1, \ell]}, \mathbf{Y}_{[1, i'-1]}) \\ &\quad + \sum_{j'=m-b+1}^m H(Y_{i', \sigma(j')} | \mathbf{Y}_{[1, i'-1]}, \tilde{\mathcal{Y}}(i', [1, i'-1], j')), \end{aligned} \quad (10)$$

where $\tilde{\mathcal{Y}}(i', [1, i'-1], j')$ is defined using (7). Using (5), each term under the first summation in (10) is upper bounded by $\min(t\alpha, (d - i' + 1)\beta)$, while each term under the second summation in (10) is upper bounded using Lemma III.1. Thus, we get that

$$\begin{aligned} &H(Y_{i', [\ell+1, m]} | Y_{i', [1, \ell]}, \mathbf{Y}_{[1, i'-1]}) \\ &\leq a \min(t\alpha, (d - i' + 1)\beta) + b \min\left(\alpha, \frac{(d - i' + 1)\beta}{t}\right) \\ &= (m - \ell) \min\left(\alpha, \frac{(d - i' + 1)\beta}{t}\right). \end{aligned} \quad (11)$$

The desired bound now follows by combining (9) with (11).

IV. GENERAL FILE SIZE BOUND, FUNCTIONAL REPAIR

In this section, we present the file-size upper bound under functional repair via IFG analysis. Under functional repair, ability to recover a file after a sequence of node failures and repairs is equivalent to multicasting the source file to an arbitrary number of data collectors over the IFG [5]. The IFG characterizes the data flows from the source to a data collector, and reflects the sequence of failures and repairs in the storage network. The IFG used here (see Fig. 3) is a generalization of the one presented in [6] for the case of $t = 1$.

A. Information Flow Graph Model

Let \mathcal{X}_i denote the physical cluster i , and let $X_{i,j}$ denote the physical node j in cluster i , $1 \leq i \leq n$, $1 \leq j \leq m$. In the IFG, $X_{i,j}$ is mapped to the pair of nodes $(X_{i,j}^{in}, X_{i,j}^{out})$ such

that the edge $X_{i,j}^{in} \rightarrow X_{i,j}^{out}$ has capacity α . The *external* node X_i^{ext} of cluster i serves to transfer data outside the cluster. The m out-nodes connect to X_i^{ext} via edges of capacity α .

When a cluster, say i , experiences a batch of t failures, the whole cluster becomes inactive and is replaced with a new active cluster. In the new cluster, a special *repair* node X_i^{rep} is used to combine local and external helper data, and generate the content of the replacement nodes. The out nodes of the ℓ local helper nodes connect to X_i^{rep} via links of capacity α , and the external nodes of the d helper clusters connect to X_i^{rep} via links of capacity β . Also, X_i^{rep} connects to the in-nodes of the replacement nodes via links of capacity α . Further, the $m - t$ nodes, which did not experience failure in the inactive cluster are copied as such in the new active cluster. At any point in time, physical cluster i contains one active and f_i inactive clusters in the IFG where $f_i \geq 0$ denotes the total number of batch failures and repairs in the cluster. We write $\mathcal{X}_i(\tau), 0 \leq \tau \leq f_i$ to denote the cluster in the IFG after the τ^{th} (batch) repair associated with cluster i , and use $\mathcal{R}_i(\tau), 0 \leq \tau \leq f_i - 1$ to denote the indices of nodes that fail in $\mathcal{X}_i(\tau)$. The clusters $\mathcal{X}_i(0), \dots, \mathcal{X}_i(f_i - 1)$ are inactive, while $\mathcal{X}_i(f_i)$ is active, after f_i repairs. The nodes of $\mathcal{X}_i(\tau)$ will be denoted by $X_{i,j}^{in}(\tau), X_{i,j}^{out}(\tau), X_i^{ext}(\tau), X_i^{rep}(\tau)$ (there is no repair node if $\tau = 0$).

Finally, the source node S connects to all the mn in-nodes $X_{i,j}^{in}(0), 1 \leq i \leq m, 1 \leq j \leq n$ via links of infinite capacity. The data collector T connects to k external nodes, say $X_i^{ext}(f_i), 1 \leq i \leq k$ also via links of infinite capacity.

B. File-Size Upper Bound

We explain the proof of the bound in (1) by considering the special case $(n = 3, k = 2, d = 2)(\alpha, \beta)(m = 3, \ell = 0, t = 2)$. A full proof appears in Appendix B. Note that for this special case, $t \nmid (m - \ell)$ and this will help us illustrate the difference between functional and exact repair. Consider the following sequence of 4 batches of failures and repairs (see Fig. 3). Batches 1 and 2 are associated with cluster 1 with $\mathcal{R}_1(0) = \{2, 3\}$ and $\mathcal{R}_1(1) = \{1, 2\}$. Batches 3 and 4 are associated with cluster 3 with $\mathcal{R}_3(0) = \{2, 3\}$ and $\mathcal{R}_3(1) = \{1, 2\}$. There is no local help in this example, cluster 1 receives external help from $X_2^{ext}(0)$ and $X_3^{ext}(0)$ for both batches of repairs, while cluster 3 receives external help from $X_2^{ext}(0)$ and $X_1^{ext}(2)$ for its repairs. Consider data collection by connecting to $X_1^{ext}(2)$ and $X_3^{ext}(2)$, and consider the S - T cut whose edges are found as follows: For disconnecting $X_{1,1}^{out}(2)$ and $X_{1,2}^{out}(2)$, we either remove (based on whichever has smaller capacity) the two edges $X_{1,1}^{in}(2) \rightarrow X_{1,1}^{out}(2)$ and $X_{1,2}^{in}(2) \rightarrow X_{1,2}^{out}(2)$ or the set of helper edges $X_2^{ext}(0) \rightarrow X_1^{rep}(2)$ and $X_3^{ext}(0) \rightarrow X_1^{rep}(2)$. For disconnecting $X_{1,3}^{out}(2)$, we either remove the *single* edge $X_{1,3}^{in}(1) \rightarrow X_{1,3}^{out}(1)$ or the set of *two* helper edges $X_2^{ext}(0) \rightarrow X_1^{rep}(1)$ and $X_3^{ext}(0) \rightarrow X_1^{rep}(1)$. The set of edges that disconnects cluster 3 is similarly found, except that if we choose to disconnect links from external helpers, we only disconnect those from $X_2^{ext}(0)$ and not $X_1^{ext}(2)$. The value of the cut forms an upper bound for B , and is given by $B \leq \min(2\alpha, d\beta) + \min(\alpha, d\beta) + \min(2\alpha, (d - 1)\beta) + \min(\alpha, (d - 1)\beta)$, which is the same as the one give by (1).

Converse: We note that it can also be shown that any valid IFG, regardless of the specific sequence of failures and repairs, B_F^* (see (1)) is indeed a lower bound on the minimum possible value of any S - T cut. Please see Appendix C for a proof of this fact, which establishes the system capacity under functional repair. Note that the RLNC simulation in Fig. 2b is an experimental verification for the validity of this converse statement.

REFERENCES

- [1] H. Abu-Libdeh, L. Princehouse, and H. Weatherspoon, "Racs: a case for cloud storage diversity," in *Proceedings of the 1st ACM symposium on Cloud computing*. ACM, 2010, pp. 229–240.
- [2] A. Bessani, M. Correia, B. Quaresma, F. André, and P. Sousa, "Depsky: dependable and secure storage in a cloud-of-clouds," *ACM Transactions on Storage (TOS)*, vol. 9, no. 4, p. 12, 2013.
- [3] J. Y. Chung, C. Joe-Wong, S. Ha, J. W.-K. Hong, and M. Chiang, "Cyrus: Towards client-defined cloud storage," in *Proceedings of the Tenth European Conference on Computer Systems*. ACM, 2015, p. 17.
- [4] J. D. Cook, R. Primmer, and A. de Kwant, "Compare cost and performance of replication and erasure coding," *Hitachi Review*, vol. 63, p. 304, 2014.
- [5] A. G. Dimakis, P. B. Godfrey, Y. Wu, M. J. Wainwright, and K. Ramchandran, "Network coding for distributed storage systems," *Information Theory, IEEE Transactions on*, vol. 56, no. 9, pp. 4539–4551, 2010.
- [6] N. Prakash, V. Abdrashitov, and M. Médard, "The Storage vs Repair-Bandwidth Trade-off for Clustered Storage Systems," *ArXiv e-prints*, vol. abs/1701.04909, Jan. 2017.
- [7] M. Gerami, M. Xiao, and M. Skoglund, "Two-layer coding in distributed storage systems with partial node failure/repair," *IEEE Communications Letters*, vol. PP, no. 99, 2017.
- [8] D. Ford, F. Labelle, F. I. Popovici, M. Stokely, V.-A. Truong, L. Barroso, C. Grimes, and S. Quinlan, "Availability in globally distributed storage systems," in *OSDI*, vol. 10, 2010, pp. 1–7.
- [9] R. Koetter and M. Médard, "An algebraic approach to network coding," *IEEE/ACM Transactions on Networking*, vol. 11, no. 5, pp. 782–795, Oct 2003.
- [10] T. Ho, M. Médard, R. Koetter, D. R. Karger, M. Effros, J. Shi, and B. Leong, "A random linear network coding approach to multicast," *IEEE Transactions on Information Theory*, vol. 52, no. 10, pp. 4413–4430, Oct 2006.
- [11] K. W. Shum and Y. Hu, "Cooperative regenerating codes," *IEEE Transactions on Information Theory*, vol. 59, no. 11, 2013.
- [12] A.-M. Kermarrec, N. Le Scouarnec, and G. Straub, "Repairing multiple failures with coordinated and adaptive regenerating codes," in *Network Coding (NetCod), 2011 International Symposium on*. IEEE, 2011.
- [13] V. R. Cadambe, S. A. Jafar, H. Maleki, K. Ramchandran, and C. Suh, "Asymptotic interference alignment for optimal repair of mds codes in distributed storage," *IEEE Transactions on Information Theory*, vol. 59, no. 5, pp. 2974–2987, 2013.
- [14] A. S. Rawat, O. O. Koyluoglu, and S. Vishwanath, "Centralized repair of multiple node failures with applications to communication efficient secret sharing," *CoRR*, vol. abs/1603.04822, 2016.
- [15] Y. Hu, P. P.-C. Lee, and X. Zhang, "Double regenerating codes for hierarchical data centers," in *Information Theory (ISIT), 2015 IEEE International Symposium on*. IEEE, 2016.
- [16] J. Sohn, B. Choi, S. W. Yoon, and J. Moon, "Capacity of clustered distributed storage," *CoRR*, vol. abs/1610.04498, 2016.
- [17] B. Gastón, J. Pujol, and M. Villanueva, "A realistic distributed storage system: the rack model," *arXiv preprint arXiv:1302.5657*, 2013.
- [18] J. Pernas, C. Yuen, B. Gastón, and J. Pujol, "Non-homogeneous two-rack model for distributed storage systems," in *Information Theory Proceedings (ISIT), 2013 IEEE International Symposium on*, July 2013.
- [19] G. Calis and O. O. Koyluoglu, "Architecture-aware coding for distributed storage: Repairable block failure resilient codes," *CoRR*, vol. abs/1605.04989, 2016.
- [20] N. B. Shah, K. V. Rashmi, P. V. Kumar, and K. Ramchandran, "Distributed storage codes with repair-by-transfer and nonachievability of interior points on the storage-bandwidth tradeoff," *IEEE Transactions on Information Theory*, vol. 58, no. 3, pp. 1837–1852, March 2012.

APPENDIX A
PROOF OF LEMMA III.1

We will show that the permutation σ_{i', S_i} is such that

$$H\left(Y_{i', \sigma_{i', S_i}(j')} | \mathcal{Y}(S_i), \tilde{\mathcal{Y}}(i', S_i, j')\right) \leq \min\left(\alpha, \frac{(d-i)\beta}{t}\right), \quad (12)$$

for all $j' \in \{m-b+1, m-b+2, \dots, m\}$, where

$$\tilde{\mathcal{Y}}(i', S_i, j') = Y_{i', [1, \ell]} \cup \{Y_{i', \sigma_{i', S_i}(j)}, j \in [\ell+1, j'-1]\}. \quad (13)$$

Consider the variable j' appearing in (12), and let $j' = m-x$ for some $x, 0 \leq x \leq b-1$ so that using (8) we have, $\sigma_{i', S_i}(j') = j_{m-x}$. Consider the definition of $(j_{m-x}, \mathcal{V}_{m-x})$ in (8); we then know that

$$H\left(Y_{i', j_{m-x}} | \mathcal{V}_{m-x}, \mathcal{Y}(S_i), Y_{i', [1: \ell]}\right) \leq H\left(Y_{i', j_p} | \mathcal{V}, \mathcal{Y}(S_i), Y_{i', [1: \ell]}\right), \quad (14)$$

for all $\mathcal{V} \subset \{Y_{i', j_{\ell+1}}, Y_{i', j_{\ell+2}}, \dots, Y_{i', j_{m-x}}\} \setminus \{Y_{i', j_p}\}$ such that $|\mathcal{V}| = t-1$, and for all $p, \ell+1 \leq p \leq m-x-1$. Towards proving (12), first of all, observe that

$$H\left(Y_{i', \sigma_{i', S_i}(j')} | \mathcal{Y}(S_i), \tilde{\mathcal{Y}}(i', S_i, j')\right) \leq H\left(Y_{i', \sigma_{i', S_i}(j')} | \mathcal{Y}(S_i), \mathcal{V}_{m-x}, Y_{i', [1: \ell]}\right). \quad (15)$$

This follows from (13) and because of the fact that $\mathcal{V}_{m-x} \subset \{Y_{i', j_{\ell+1}}, Y_{i', j_{\ell+2}}, \dots, Y_{i', j_{m-x-1}}\}$. Without loss of generality, assume that $\mathcal{V}_{m-x} = \{Y_{i', j_{\ell+1}}, Y_{i', j_{\ell+2}}, \dots, Y_{i', j_{\ell+t-1}}\}$. Next, from the exact repair condition given in (5), we know that

$$\begin{aligned} \min(t\alpha, (d-i)\beta) &\geq H\left(Y_{i', \sigma_{i', S_i}(j')}, \mathcal{V}_{m-x} | \mathcal{Y}(S_i), Y_{i', [1: \ell]}\right) \\ &= \sum_{p=\ell+1}^{\ell+t-1} H\left(Y_{i', j_p} | Y_{i', j_{\ell+1}}, \dots, Y_{i', j_{p-1}}, \mathcal{Y}(S_i), Y_{i', [1: \ell]}\right) + \\ &\quad H\left(Y_{i', \sigma_{i', S_i}(j')} | \mathcal{V}_{m-x}, \mathcal{Y}(S_i), Y_{i', [1: \ell]}\right) \\ &\geq \sum_{p=\ell+1}^{\ell+t-1} H\left(Y_{i', j_p} | \mathcal{V}_{j_p}, \mathcal{Y}(S_i), Y_{i', [1: \ell]}\right) + \\ &\quad H\left(Y_{i', \sigma_{i', S_i}(j')} | \mathcal{V}_{m-x}, \mathcal{Y}(S_i), Y_{i', [1: \ell]}\right), \quad (16) \end{aligned}$$

where $\mathcal{V}_{j_p} = \mathcal{V}_{m-x} \setminus \{Y_{i', j_p}\} \cup \{Y_{i', \sigma_{i', S_i}(j')}\}$. Noting that $|\mathcal{V}_{j_p}| = t-1$, we see that each term under the first summation in (16) can be lower bounded using (14), i.e.,

$$\begin{aligned} H\left(Y_{i', j_p} | \mathcal{V}_{j_p}, \mathcal{Y}(S_i), Y_{i', [1: \ell]}\right) &\geq H\left(Y_{i', j_{m-x}} | \mathcal{V}_{m-x}, \mathcal{Y}(S_i), Y_{i', [1: \ell]}\right) \\ &= H\left(Y_{i', \sigma_{i', S_i}(j')} | \mathcal{V}_{m-x}, \mathcal{Y}(S_i), Y_{i', [1: \ell]}\right) \end{aligned} \quad (17)$$

Combining (17) with (16), it follows that

$$H\left(Y_{i', \sigma_{i', S_i}(j')} | \mathcal{V}_{m-x}, \mathcal{Y}(S_i), Y_{i', [1: \ell]}\right) \leq \min\left(\alpha, \frac{(d-i)\beta}{t}\right). \quad (18)$$

The proof of lemma now follows by combining (18) with (15).

APPENDIX B
PROOF OF UPPER-BOUND 1

We prove that under functional repair the file-size is upper-bounded by

$$\begin{aligned} B &\leq B_F^* = lk\alpha + a \sum_{i=1}^k \min(t\alpha, (d-i+1)\beta) \\ &\quad + \sum_{i=1}^k \min(b\alpha, (d-i+1)\beta). \end{aligned}$$

Let $[A, B] = \{\text{integer } x : A \leq x \leq B\}$, and $[B] = [1, B]$.

To show the bound, it is enough to demonstrate a sequence of batch failures and a set of k clusters used by a data-collector, such that there exists a cut between the source and the data-collector with capacity no more than B_F^* . In the example sequence, that we consider, the clusters 1 to k are used for data-collection and experience node failures. At each of these clusters $a+1$ batch failures occur. They jointly cover the first $m-\ell$ nodes of a cluster. Specifically, at cluster $i \in [k]$, the first batch failure affects the last t of these nodes: $\mathcal{R}_i(0) = \{m-\ell-t+1, \dots, m-\ell\}$. The remaining batch failures affect disjoint sets of t nodes starting from the first node $X_{i,1}$: $\mathcal{R}_i(1) = \{1, \dots, t\}$, $\mathcal{R}_i(2) = \{t+1, \dots, 2t\}$, until $\mathcal{R}_i(a) = \{(a-1)t, \dots, at\}$.

In all cases, the last ℓ nodes in a cluster provide the local helper data. For repairs in cluster i , clusters $1, \dots, i-1$ and $n-(d-i), \dots, n$ serve as helper clusters. Failures first occur in cluster 1, then in clusters 2, 3, etc. until cluster k .

In the IFG, corresponding to the described failure sequence, cluster $\mathcal{X}_i(a+1)$ is active for each $i \in [k]$. Let τ_j be the such that the cluster $\mathcal{X}_i(\tau_j)$ appears in the IFG right after the last repair of node $X_{i,j}$ (we say ‘‘last repair’’ since nodes whose indices belong to $\mathcal{R}_i(0) \cap \mathcal{R}_i(a)$ fail twice in our sequence of failures; other nodes in cluster i fail only once). Consider a cut-set (IFG_S, IFG_T) consisting of the following edges:

- $X_{i,j}^{in}(a+1) \xrightarrow{\alpha} X_{i,j}^{out}(a+1), \forall i \in [k], j \in [m-\ell+1, m]$. Total capacity of these edges is $lk\alpha$.
- For all $i \in [k]$:
 - Edge set $X_i^{rep}(\tau_j) \xrightarrow{\alpha} X_{i,j}^{in}(\tau_j), j \in [at]$, or edge set $X_i^{ext}(0) \xrightarrow{\beta} X_i^{rep}(\tau_j) \forall i' \in [n-(d-i), n], j \in \{t, 2t, \dots, at\}$, whichever set capacity is smaller. Total capacity of these edges is $a \min(t\alpha, (d-i+1)\beta)$.
 - If $b > 0$: edge set $X_i^{rep}(\tau_j) \xrightarrow{\alpha} X_{i,j}^{in}(\tau_j), j \in [at+1, m-\ell]$, or edge set $X_i^{ext}(0) \xrightarrow{\beta} X_i^{rep}(\tau_j) \forall i' \in [n-(d-i), n], j = m-\ell$, whichever set capacity is smaller. Total capacity of these edges is $\min(b\alpha, (d-i+1)\beta)$.

The value of the cut is given by $lk\alpha + a \sum_{i=1}^k \min(t\alpha, (d-i+1)\beta) + \sum_{i=1}^k \min(b\alpha, (d-i+1)\beta) = B_F^*$. \square

APPENDIX C

MIN-CUT FOR IFG, OPTIMALITY OF B_F^* FOR GENERAL ℓ

We now show that for any valid IFG, regardless of the specific sequence of failures and repairs, B_F^* is indeed a lower bound on the minimum possible value of any S - T cut.

Consider a cut of IFG, and let IFG_S and IFG_T be the two disjoint parts associated with nodes S and T , respectively. Without loss of generality, we only consider cuts such that IFG_T contains at least k external nodes corresponding to active clusters. Consider a topological sorting of the IFG nodes such that: 1) an edge exists between two nodes A and B only if A appears before B in the sorting, and 2) all in-, out-, external, and repair nodes (if $\tau > 0$) of the cluster $\mathcal{X}_i(\tau)$ appear together in the sorted order, $\forall i, \tau$.

Consider the sequence \mathcal{E} of all the external nodes in both active and inactive clusters in IFG_T in their sorted order. Let Y_1 denote the first node in \mathcal{E} . Without loss of generality let $Y_1 = X_1^{\text{ext}}(\tau_1)$, for some τ_1 . In this case, consider the subsequence of \mathcal{E} which is obtained after excluding all the external nodes associated with \mathcal{X}_1 from \mathcal{E} . Let Y_2 denote the first external node in this subsequence. We continue in this manner until we find the first k external nodes $\{Y_1, Y_2, \dots, Y_k\}$ in \mathcal{E} , such that each of the k nodes corresponds to a distinct physical cluster. Without loss of generality, let us also assume that $Y_i = X_i^{\text{ext}}(\tau_i)$, $2 \leq i \leq k$, for some τ_i . If $\tau_i = 0$, then clearly cluster i contributes (at least) $m\alpha$ to the cut. Thus let us assume that $\tau_i > 0$, $1 \leq i \leq k$.

Consider the m out-nodes $X_{i,1}^{\text{out}}(\tau_i), \dots, X_{i,m}^{\text{out}}(\tau_i)$ that connect to $X_i^{\text{ext}}(\tau_i)$. For each $j \in [1, m]$, either $X_{i,j}^{\text{out}}(\tau_i)$ is in IFG_S or there exists a minimal $\tau_{i,j} \in [0, \tau_i]$ such that $X_{i,j}^{\text{out}}(\tau_{i,j}) \in \text{IFG}_T$. Consider those values of $j \in [1, m]$ for which all the following conditions hold:

$$\begin{aligned} X_{i,j}^{\text{out}}(\tau_i), X_{i,j}^{\text{in}}(\tau_{i,j}) &\in \text{IFG}_T, j \in \mathcal{R}_i(\tau_{i,j} - 1), \\ X_i^{\text{rep}}(\tau_{i,j}) &\in \text{IFG}_T. \end{aligned} \quad (19)$$

Let there be $m_i \in [0, m]$ of such values, and, without loss of generality, let them be $m - m_i + 1, \dots, m$. Also without loss of generality, let indices j be sorted in the order of increasing $\tau_{i,j}$, i.e. $j_1 < j_2$ implies $\tau_{i,j_1} \leq \tau_{i,j_2}$. For each $j \in [m - m_i + 1, m]$, $\Sigma_{i,j} \triangleq \{j' : \tau_{i,j'} = \tau_{i,j}, j' \in [m - m_i + 1, m]\}$ is a contiguous set of at most t indices of the nodes with the same $\tau_{i,j}$, and which are repaired together from the same repair node. Let $\mathcal{S}_i = \{\text{distinct}(\min \Sigma_{i,j} - 1), \forall j \in [m - m_i + 1, m]\} \subseteq [m - m_i, m - 1]$ be the set of indices of the nodes preceding all contiguous groups $\Sigma_{i,j}$. Note that by $\min \Sigma_{i,j}$ we mean the minimum element contained in the set $\Sigma_{i,j}$. The set \mathcal{S}_i is in one-to-one correspondence with the set of the repair nodes in (19) for $j \in [m - m_i + 1, m]$. Note that $m - m_i$ is always an element of \mathcal{S}_i .

In order to relay helper data to $X_{i,j}^{\text{in}}(\tau_{i,j})$ for all $j \in [m - m_i + 1, m]$, the number of these repair nodes should be at least $\lceil m_i/t \rceil$, and $|\mathcal{S}_i| \geq \lceil m_i/t \rceil$. Each of these repair nodes connects to d external nodes in other clusters. By construction of \mathcal{E} , at most $i - 1$ of those external nodes can be in IFG_T . Thus, each repair node contributes at least $(d - i + 1)\beta$ of external helper data to the cut value. In addition, each repair node $X_i^{\text{rep}}(\tau_{i,j})$ connects to ℓ local nodes. By (19) and by construction of \mathcal{S}_i and sorting of $\tau_{i,j}$, only nodes with indices $\{1, 2, \dots, j'\}$ out of these ℓ can be in IFG_T , where $j' = \min \Sigma_{i,j} - 1$ is the corresponding element of \mathcal{S}_i . Thus, repair node $X_i^{\text{rep}}(\tau_{i,j})$ contributes at least $(\ell - j')^+\alpha$ of local helper data to the cut value.

The contribution to the cut value of those $m - m_i$ indices of $j \in [1, m - m_i]$, which do not satisfy (19), is at least α each.

Based on the observations above, the overall cut value is lower-bounded by

$$\begin{aligned} &\text{mincut}(S - T) \\ &\geq \sum_{i=1}^k ((m - m_i)\alpha + \lceil \frac{m_i}{t} \rceil (d - i + 1)\beta + \sum_{j' \in \mathcal{S}_i} (\ell - j')^+\alpha). \end{aligned} \quad (20)$$

Consider a particular value of $i \in [1, k]$ and the corresponding summation term in (20). Let us assume that $m - m_i \geq \ell$, and $m_i = a_i t + b_i \leq m - \ell$, $b_i \in [0, t - 1]$. Then the third term in (20) is zero, and

$$\begin{aligned} &(m - m_i)\alpha + \lceil m_i/t \rceil (d - i + 1)\beta \\ &= m\alpha - (a_i t + b_i)\alpha + (a_i + 1_{b_i > 0})(d - i + 1)\beta \\ &= m\alpha - a_i(t\alpha - (d - i + 1)\beta) - b_i\alpha + 1_{b_i > 0}(d - i + 1)\beta \\ &\stackrel{(1)}{\geq} \ell\alpha + (m - \ell)\alpha - a(t\alpha - (d - i + 1)\beta)^+ \\ &\quad - (b\alpha - (d - i + 1)\beta)^+ \\ &= \ell\alpha + a(t\alpha - (t\alpha - (d - i + 1)\beta)^+) \\ &\quad + (b\alpha - (b\alpha - (d - i + 1)\beta)^+) \\ &= \ell\alpha + a \min(t\alpha, (d - i + 1)\beta) + \min(b\alpha, (d - i + 1)\beta) \\ &\triangleq C_i, \end{aligned}$$

where (1) follows, because $a_i t + b_i = m_i \leq m - \ell = at + b$, $a_i \leq a$, and, if $a_i = a$, $b_i \leq b$.

On the other hand, if $m - m_i = \ell - \mu_i < \ell$, and $m_i > m - \ell = at + b$, $\ell - (m - m_i) = \mu_i > 0$, then we have

$$\begin{aligned} &(m - m_i)\alpha + \lceil m_i/t \rceil (d - i + 1)\beta + \sum_{j' \in \mathcal{S}_i} (\ell - j')^+\alpha \\ &\geq (\ell - \mu_i)\alpha + (a + 1_{b > 0})(d - i + 1)\beta + (\ell - (m - m_i))\alpha \\ &\quad + \sum_{\substack{j' \in \mathcal{S}_i \\ j' > m - m_i}} (\ell - j')^+\alpha \\ &= \ell\alpha + (a + 1_{b > 0})(d - i + 1)\beta + \sum_{\substack{j' \in \mathcal{S}_i \\ j' > m - m_i}} (\ell - j')^+\alpha \\ &\geq C_i, \end{aligned}$$

where C_i is the lower-bound for the case $m - m_i \geq \ell$.

Since $B_F^* = \sum_i C_i$, it is indeed a lower-bound on the file-size. This proves tightness of 1.