# Balanced and Swap-Robust Trades for Dynamical Distributed Storage 

Chao Pan<br>ECE Department<br>University of Illinois<br>chaopan2@illinois.edu

Ryan Gabrys<br>Calit2<br>University of California-San Diego<br>ryan.gabrys@gmail.com

Xujun Liu<br>Foundational Mathematics Department<br>Xi'an Jiaotong-Liverpool University<br>xujun.liu@xjtlu.edu.cn

Charles Colbourn<br>School of CAI<br>Arizona State University<br>charles.colbourn@asu.edu

Olgica Milenkovic<br>ECE Department<br>University of Illinois<br>milenkov@illinois.edu


#### Abstract

Trades, introduced by Hedayat [9], are two sets of blocks of elements which may be exchanged (traded) without altering the counts of certain subcollections of elements within their constituent blocks. They are of importance in applications where certain combinations of elements dynamically become prohibited from being placed in the same group of elements, since in this case one can trade the offending blocks with allowed ones. This is particularly the case in distributed storage systems, where due to privacy and other constraints, data of some groups of users cannot be stored together on the same server. We introduce a new class of balanced trades, important for access balancing of servers, and perturbation resilient balanced trades, important for studying the stability of server access frequencies with respect to changes in data popularity. The constructions and bounds on our new trade schemes rely on specialized selections of defining sets in minimal trades and number-theoretic analyses.


Index Terms-access balancing, combinatorial trades, distributed storage, data popularity changes

## I. Introduction

An important problem in distributed storage is to balance access requests to servers in order to prevent service time bottlenecks. Access balancing aims to limit the number of server session logon requests and queues excess requests until server resources become available. As data access directly relates to the popularity of items on a server, it is important to design storage systems that balance the popularity of data chunks on each server.

Existing distributed storage systems categorize data into hot, warm and cold based on their popularity and use different mechanisms for storing them [4], [11], [14]. High-level partitions like these fail to account for popularity-variability within each class and the resulting dynamic changes in data demands, and they do not couple the balancing process with distributed storage repair solutions [7], [8], [15]. As a result, although issues such as delay-storage tradeoffs, volume (load) balancing, and chunk allocation have been studied in depth [1], [11], [12], access balancing has received significantly less attention. To

[^0]address this problem, the authors of [6] introduced MaxMin Steiner triple systems and related combinatorial designs for fractional repetition distributed storage coding [8]. MaxMin Steiner triple systems maximize the minimum average popularity score of items stored on different servers in the system and thereby ensure control of the discrepancy of average server popularity scores [2], [3], [5]. These systems are constructed by selecting a specific packing or design, and associating data items of given popularities with elements of the packing or design. The MaxMin Steiner balancing approach can also be integrated with batch codes [8], [15], [16].

In [9], the following question was posed: "In the application of $t$-designs we may be confronted with a situation where some blocks become too costly to be selected for experimentation. For example, combining certain tasks within one block may be unacceptable to the experimenter. Suppose the available $t$-designs in the literature contain such blocks and by renaming we cannot dispose of these undesirable blocks. The theory of trade-offs can tell us how to trade undesirable blocks with those which are acceptable to the experimenter." In access balancing such situations arises when one needs to redistribute files or data chunks on servers to account for constraints on items which cannot be stored together or to mitigate the effect of dynamical changes in the popularity of items.

This work demonstrates that such constrained trading problems may be resolved through the use of new classes of balanced and swap-robust trades. A trade with parameters $(v, k, t)$ is a pair of disjoint sets of blocks $T^{(1)}, T^{(2)}$ of equal size over a ground set of $v$ points such that all blocks are of size $k$ and any subset of $t$ points appears the same number of times in blocks of $T^{(1)}$ and blocks of $T^{(2)}$. The set of blocks in $T^{(1)}$ and $T^{(2)}$ are subgroups of blocks in combinatorial designs which are used for distributed and batch code constructions. Replacing blocks in $T^{(1)}$ by blocks in $T^{(2)}$ when the need arises does not change the global block intersection properties. In the setting of access-balanced storage with dynamically changing item popularities, undesirable blocks may represent blocks whose average popularity has decreased or increased significantly or
blocks who are dynamically faced with storing disallowed combinations of data items. Importantly, for dynamic access balancing, blocks in $T^{(1)}$ can be traded by blocks in $T^{(2)}$ without changing the properties of the underlying design but changing the points (e.g., data chunks or files) allocated to the blocks (e.g., servers). Due to space limitations, we relegate the detailed description of how to use trades in distributed storage system to the journal version of this work.

The first question of interest is: Can trades themselves be made "balanced"? Labeling the points with numbers in $\{1,2, \ldots, v\}$, the question of balancing translates to ensuring that all blocks in $T^{(1)}$ and $T^{(2)}$ have the same sum of label values. Our first result is that minimal trades (i.e., trades with the smallest number of blocks) can be easily balanced (see Section III). Balancing is needed even when no changes in popularities are present to ensure maintaining near-uniform server access; minimal trades are of interest since they require replacing (trading) the smallest number of blocks to achieve the desired rebalancing or to meet a given placement constraint. The second question is how to design balanced trades that are as robust as possible to limited-magnitude swaps in popularities of items? For $v=4(t+1)$, where $t$ is a positive integer, we show that there exist balanced trades whose popularity sums do not change more than roughly $2(t+1)$ when all popularity values are allowed to undergo arbitrary adjacent swaps. Interestingly, for $v=16$ and $v=20$ the corresponding balanced minimal trades are unique.

## II. Basics of Trades

We start our exposition with some important definitions.
Given a set of elements $V=[v]=\{1, \ldots, v\}$, let $P(v)$ denote the power set of $V$ and let $P_{k}(v) \subseteq P(v)$ be the set of all subsets (blocks) of $P(v)$ that have cardinality $k$, for some positive integer $k<v$.
Definition 1. $A(v, k, t)$ trade is a pair of sets $\left\{T^{(1)}, T^{(2)}\right\}$, $T^{(i)} \subseteq P_{k}(v), i=1,2$, such that $\left|T^{(1)}\right|=\left|T^{(2)}\right|$ and $T^{(1)} \cap T^{(2)}=\varnothing$, satisfying the following property: For any $B_{t} \in P_{t}(v)$, the number of blocks in $T^{(1)}$ that contain $B_{t}$ is the same as the number of blocks in $T^{(2)}$ that contain $B_{t}$. The volume of the trade is the number of blocks in $T^{(1)}$ and $\left|T^{(1)}\right|=\left|T^{(2)}\right| . A(v, k, t)$ trade is termed minimal if it has the smallest possible number of blocks.

We start by describing a construction of minimal trades based on Hwang [10] which allows one to balance the sum of elements in blocks (i.e., ensure discrepancy $=0$ ) of both $T^{(1)}$ and $T^{(2)}$, using specialized choices of the sets defining the blocks of the trades. We describe the construction and an accompanying new proof that immediately reveals how to perform balancing (note that no proof for this result was provided in [10]). For each $\ell \in\{1,2, \ldots, n\}$, let
$P_{\ell}=\left\{\left(i_{1}, i_{1}+1\right) \cdots\left(i_{\ell}, i_{\ell}+1\right) \mid\left\{i_{1}, \ldots, i_{\ell}\right\} \subseteq\{1, \ldots, 2 n-1\}\right\}$,
be a set of permutations over a set of $2 n$ elements, where $\left(i_{1}, i_{1}+1\right)\left(i_{2}, i_{2}+1\right) \ldots\left(i_{\ell}, i_{\ell}+1\right)$ denotes the permutation which swaps $i_{j}$ with $i_{j}+1$, for $j=1, \ldots, \ell$, and leaves all
other elements fixed. Furthermore, let $P_{0}=\{(e)\}$ consist of the identity permutation $e$. Also, define

$$
\Delta_{2 n}=\bigcup_{\ell \text { even }} P_{\ell} \quad \text { and } \quad \bar{\Delta}_{2 n}=\bigcup_{\ell \text { odd }} P_{\ell}
$$

Theorem 2. ( 10$\rceil$ ) There exists a $(v, k, t)$ trade of volume $2^{t}$. This volume is the smallest (minimal) possible volume for the given choice of parameters.

Proof: We provide a simple proof for the first claim. Let $S_{1}, S_{2}, \ldots, S_{2 t+3}$ be subsets of the point set $V$ that can be partitioned into $t+1$ pairs of the form

$$
S_{1}, S_{2} ; S_{3}, S_{4} ; \ldots S_{2 t+1}, S_{2 t+2}
$$

with the addition of one unpaired set $S_{2 t+3}$. We henceforth refer to the sets as defining sets and the pairs as companions. Assume that the defining sets have the following properties:

1) $S_{i} \cap S_{j}=\varnothing$ for $i \neq j$;
2) $\left|S_{2 i-1}\right|=\left|S_{2 i}\right| \geqslant 1$, for $i=1, \ldots, t+1$;
3) $\sum_{i=1}^{t+2}\left|S_{2 i-1}\right|=k$.

Respectively, define $T=\left\{T^{(1)}, T^{(2)}\right\}$ as

$$
\begin{aligned}
& \left\{\left(S_{\sigma(1)} \cup S_{\sigma(3)} \cup \ldots \cup S_{\sigma(2 t+1)} \cup S_{2 t+3}\right): \sigma \in \Delta_{2 t+2}\right\} \\
& \left\{\left(S_{\bar{\sigma}(1)} \cup S_{\bar{\sigma}(3)} \cup \ldots \cup S_{\bar{\sigma}(2 t+1)} \cup S_{2 t+3}\right): \bar{\sigma} \in \bar{\Delta}_{2 t+2}\right\}
\end{aligned}
$$

We used $\left(S_{\sigma(1)} \cup S_{\sigma(3)} \cup \ldots \cup S_{\sigma(2 t+1)} \cup S_{2 t+3}\right)$ and $\left(S_{\bar{\sigma}(1)} \cup\right.$ $\left.S_{\bar{\sigma}(3)} \cup \ldots \cup S_{\bar{\sigma}(2 t+1)} \cup S_{2 t+3}\right)$ to denote blocks whose elements represent the unions of the corresponding sets $S$.
Claim 3. The trade $T$ has volume $2^{t}$.
Proof: Since $\Delta_{2 t+2}$ represents the set of all permutations with an even number of transpositions, for $t$ even we have
$\left|T^{(1)}\right|=\left|\Delta_{2 t+2}\right|=\binom{t+1}{0}+\binom{t+1}{2}+\ldots+\binom{t+1}{t}=2^{t}$,
$\left|T^{(2)}\right|=\left|\bar{\Delta}_{2 t+2}\right|=\binom{t+1}{1}+\binom{t+1}{3}+\ldots+\binom{t+1}{t+1}=2^{t}$. The same results hold for odd $t$.

Claim 4. $T$ is a $(v, k, t)$-trade.
Proof: We show that every block has size $k$ and that every $t$-subset of $V$ is contained in the same number of blocks in $T^{(1)}$ and $T^{(2)}$.

Firstly, the claim that every block has size $k$ follows from Property (3) of the collection of sets $S_{1}, S_{2}, \ldots, S_{2 t+3}$. Secondly, to show that every $t$-subset of $V$ is contained in the same number of blocks in $T^{(1)}$ and $T^{(2)}$, let $U$ be a $t$-subset of the ground set $V$.

Case 1: There exists some $i \in\{1, \ldots, t+1\}$ such that $U \cap$ $S_{2 i-1} \neq \varnothing$ and $U \cap S_{2 i} \neq \varnothing$. Then, the construction of the trade ensures that $U$ is not contained in any block.

Case 2: For each pair $S_{2 i-1}, S_{2 i}, 1 \leqslant i \leqslant t+1, U$ has a nonempty intersection with at most one of the sets. Without loss of generality, assume that $U \cap S_{2 i-1} \neq \varnothing$ for indices $i \in\left\{i_{1}, i_{2}, \ldots, i_{h_{1}}\right\}$ and $U \cap S_{2 j} \neq \varnothing$ for indices
$j \in\left\{k_{1}, k_{2}, \ldots, k_{h_{2}}\right\}$. Furthermore, without loss of generality, assume that $\left(h_{1}+h_{2}\right)$ is odd and that $t$ is even. Then $U$ is contained in

$$
\binom{t+1-h_{1}-h_{2}}{1}+\ldots+\binom{t+1-h_{1}-h_{2}}{t-h_{1}-h_{2}}
$$

blocks of $T^{(1)}$ and

$$
\binom{t+1-h_{1}-h_{2}}{0}+\ldots+\binom{t+1-h_{1}-h_{2}}{t+1-h_{1}-h_{2}}
$$

blocks of $T^{(2)}$. In both cases, the sums equal $2^{t-h_{1}-h_{2}}$. This establishes the claim.

This completes the proof of the theorem.
We remark that all trades of volume $2^{t}$ must have a structure as described in the above proof, as shown in [10].

## III. Balanced Minimal Trades

We introduce next the notion of block-sums and blockdiscrepancies needed for our subsequent discussion.
Definition 5. Let $B=\left(b_{1}, b_{2}, \ldots, b_{k}\right)$ be a block in $T^{(i)}$, where $i \in\{1,2\}$. The block-sum of $B$ equals $\Sigma_{B}=\sum_{i=1}^{k} b_{i}$, while the minimum and maximum block sums of $T^{(i)}$ are defined as $\min _{B \in T^{(i)}} \Sigma_{B}$ and $\max _{B \in T^{(i)}} \Sigma_{B}$, for $i \in\{1,2\}$. The block-discrepancy of $T^{(i)}$ is defined as $\max _{B \in T^{(i)}} \Sigma_{B}$ $\min _{B \in T^{(i)}} \Sigma_{B}$, for $i \in\{1,2\}$.
Henceforth, we tacitly assume that the points are labeled according to their popularity score so that no two popularities are the same: Label " 1 " indicates the most popular data chunk, while label " $v$ " denotes the least popular data chunk.

We ask the following questions: Can we ensure that both $T^{(1)}$ and $T^{(2)}$ have zero block-discrepancy? We affirmatively answer this question in the following lemma.
Lemma 6. There exists a minimal $(v, k, t)$ trade $T=$ $\left\{T^{(1)}, T^{(2)}\right\}$ of volume $2^{t}$ with $v=4(t+1)$ points such that both $T^{(1)}$ and $T^{(2)}$ are perfectly balanced: Every block has the same size $k=2(t+1)$ and block-sum equal to $(t+1)(4 t+5)$.

Proof: Let $T=\left\{T^{(1)}, T^{(2)}\right\}$ be defined as in the previous theorem, and select the companion sets as

$$
S_{1}=\{1,4\}, S_{2}=\{2,3\}, S_{3}=\{5,8\}, S_{4}=\{6,7\}, \ldots
$$

$S_{2 t+1}=\{4 t+1,4 t+4\}, S_{2 t+2}=\{4 t+2,4 t+3\}, S_{2 t+3}=\varnothing$.
Clearly, this choice of sets satisfies the conditions required of the construction to result in a trade. It is also easy to see that every block has the same size $2 t+2$ and that the sum of the labels in each block equals $(t+1)(4 t+5)$.

Another obvious choice for the defining sets $S_{1}, S_{2}, \ldots, S_{2 t+2}$ for doubly-even $v$ is
$S_{1}=\{1, v-1\}, S_{2}=\{2, v-2\} ; S_{3}=\{3, v-3\}, S_{4}=\{4, v-4\} ;$ $\ldots S_{2 t+1}=\{2 t+1, v-2 t-1\}, S_{2 t+2}=\{2 t+2, v-2 t-2\}$.
We determine next how many balanced partitions into pairs of companion sets of cardinality two, i.e. partitions such that $\sum_{\ell \in S_{2 i-1}} \ell=\sum_{\ell \in S_{2 i}} \ell$, where $i=1,2, \ldots, t+1$. To
the best of the authors' knowledge, this question has not been previously addressed in the literature and appears to be difficult. Computer search shows that for $v=12,16,20,24$ the number of valid partitions equals $86,1990,74323,4226026$. A straightforward, yet loose, lower bound, equals $\frac{(2(t+1))!}{(t+1)!2^{t+1}}$. This follows as one can first choose a pair of integers in $\{1,2, \ldots, 2(t+1)\}$, say $a$ and $b$, and then select $a+2(t+1)$ and $b+2(t+1)$ in $\{2(t+1)+1,2(t+1)+2, \ldots, 4(t+1)\}$ as these four elements constitute two balanced companion sets.

An upper bound can be obtained by viewing the problem of counting balanced set partitions as a problem pertaining to integer partitions. Specifically, the sum of all elements in $\{1,2, \ldots, v\}$ equals $s=\frac{v(v+1)}{2}=2(t+1)(4 t+5)$ and needs to be partitioned into $2(t+1)$ positive parts. Each part represents the sum of two elements in one defining set. A result from [13] shows that the number of partitions of an integer with additional constraints on the number and size of the parts can be determined asymptotically as follows. Let $V_{i}$ denote the number of parts of size $i$ in a randomly chosen partition, where $i=1,2, \ldots$. Assume that $V_{1}, V_{2}, \ldots$ are independent and geometrically distributed random variables with parameters $\lambda, \lambda^{2}, \ldots$, respectively, where $\lambda=\exp \left(-\pi \sqrt{\frac{1}{6 n}}\right)$. Then, the asymptotic partition count of interest can be obtained by performing all computations under the previous model. The formula for $\lambda$ is a consequence of the well-known asymptotic formula by Hardy and Ramanujan for the number of partitions of a positive integer $n, p(n) \sim \frac{1}{4 \sqrt{3} n} \exp \left(\pi \sqrt{\frac{2 n}{3}}\right)$.

Note that by the definition of balanced defining sets, the smallest possible sum of a defining set is 5 (i.e., $1+4$ ), and the largest possible sum is $2 v-3$ (i.e., $v-3+v$ ), implying that $V_{i} \neq 0$ and has to be even for $5 \leqslant i \leqslant 2 v-3$. Therefore, the overall probability of drawing valid constrained partition for an integer that is the sum $s$ equals

$$
\begin{aligned}
P_{s} & =\sum_{x_{5}, \ldots, x_{2 v-3}} \lambda^{2\left[5 x_{5}+6 x_{6}+\cdots+(2 v-3) x_{2 v-3}\right]} \prod_{i=1}^{s}\left(1-\lambda^{i}\right) \\
& =\binom{\frac{v}{4}+2 v-8}{2 v-8} \lambda^{s} \prod_{i=1}^{s}\left(1-\lambda^{i}\right) ; \lambda=\exp \left(-\pi \sqrt{\frac{1}{6 s}}\right)
\end{aligned}
$$

The last line follows from the facts that $\sum_{j=5}^{2 v-3} j x_{j}=\frac{s}{2}$ and $\sum_{j=5}^{2 v-3} x_{j}=\frac{v}{4}$. For the sum of two elements in each defining set, there are at most $\frac{v}{2}$ possibilities. Therefore, the final asymptotic upper bound for the number of valid partitions equals $P_{s} \times p(s) \times \frac{(v / 2)^{v / 2}}{(v / 2)!}$. Details of the proof are omitted.

## IV. Balanced and Swap-Robust Minimal Trades

Among the large number of valid balanced defining sets, .we now wish to select those that are most resilient to what we call popularity-swaps. We introduce the following definitions.
Definition 7. We say that a data chunk labeled $i$, where $i \in$ [ $v]$, experienced a popularity change of magnitude $p$ if its label changes to a value $i_{p}, i_{p} \in[v]$, so that $\left|i-i_{p}\right| \leqslant p$. Note that the most and least popular items, labeled by 1 and $v$, can only
increase or decrease their popularity, respectively and hence the rankings are not to be viewed as cyclic. Since in the presence of popularity changes the structure of a ranking (permutation) needs to be preserved, a popularity increase has to be matched by a popularity decrease. We say that $(i, j), i, j \in[v]$, experience a popularity change of magnitude $p$ if the rankings of $i$ and $j$ swap and $|i-j|=p$.
For example, the transposition $(1,2)$ corresponds to the topranked element becoming the second-ranked one and vice versa. The magnitude of the popularity change equals $p=1$. As another example, for $p=1$, and $v=4$, allowed popularity swaps include $(1,2),(2,3),(3,4)$, but only collections of swaps that do not involve the same element (number) can be used simultaneously. Two "adjacent" swaps involving one common element are considered as one "non-adjacent" swap that could result in larger popularity drops by our definition. For example, allowing both swaps $(1,2)$ and $(2,3)$ simultaneously would result in 1 dropping to 3 and 3 rising to 1 , which is a popularity change of magnitude $p=2$. For simplicity, we focus on the case $p=1$ and only provide a sample result for $p=2$.

Example 8. Suppose that $t=3$ and that the defining sets are defined as in the proof of Lemma 6 Based on Theorem 22, the following two blocks belong to $T^{(T)}$ :

$$
B_{1}=\{1,4,6,7,9,12,14,15\}, B_{2}=\{2,3,5,8,10,11,13,16\}
$$

Consider the swaps $(1,2),(4,5),(7,8),(9,10),(12,13)$, $(15,16)$ which change the defining sets $S$ to
$S_{1}=\{2,5\}, S_{2}=\{1,3\} ; S_{3}=\{4,7\}, \quad S_{4}=\{6,8\} ;$
$S_{5}=\{10,13\}, \quad S_{6}=\{9,11\} ; S_{7}=\{12,15\}, \quad S_{8}=\{14,16\}$,
and the blocks $B_{1}$ and $B_{2}$ to

$$
B_{1}=\{2,5,6,8,10,13,14,16\}, B_{2}=\{1,3,4,7,9,11,12,15\} .
$$

The discrepancy of the sum of the two blocks equals 12, which equals $3(t+1)$. In general, the block-discrepancy can increase to $3(t+1)$ if each pair of sets $S_{2 i-1}, S_{i}$ is modified by increasing/decreasing the values of two elements from one of the sets $S \in\left\{S_{2 i-1}, S_{i}\right\}$ and decreasing/increasing the values of two elements $\left\{S_{2 i-1}, S_{i}\right\} \backslash S$; this results in a contribution of 3 to the total change in the block sum for every pair of defining sets. Assuming each of the defining sets $S$ contains exactly two elements and a similar reasoning as described above, it is clear that the block-discrepancy can increase to at most $4(t+1)$ following popularity changes of magnitude $p=1$.
Based on the example, it appears that different choices of balanced defining sets of minimal trades may have different tolerance to popularity swaps of magnitude $p=1$. The question of interest is to determine which choices of balanced defining sets have the smallest sums of discrepancies between companion sets, termed total set discrepancy, under worst-case popularity changes of magnitude $p=1$.
Lemma 9. The total set discrepancy of minimal trades with $(v=4(t+1), k=2(t+1), t)$ and defining sets of cardinality

2 , under popularity changes of magnitude $p=1$, is at least $\frac{2}{3}\left(t+\frac{2}{3}\right)$.

Proof: We model the set of potential popularity changes of magnitude 1 as edges of a digraph $G=(V, E)$ with $|V|=$ $2(t+1)$. The graph $G$ has the following properties:

1) Each vertex $v_{i} \in V$ corresponds to the defining set $S_{i}$;
2) There is a directed arrow from vertex $v_{i_{1}}$ to $v_{i_{2}}$, denoted by $\left(v_{i_{1}}, v_{i_{2}}\right)$, if $\exists s$ such that $s \in S_{i_{1}}$ and $s+1 \in S_{i_{2}}$. We allow for parallel edges but do not allow self-loops.
Since the trade is balanced, if any defining set $S_{i}$ contains a pair of consecutive elements, then the elements in its companion set cannot contain a pair of consecutive integers. Since self-loops are not allowed, it follows that the number of edges in $E$ is at least $3(t+1)-1$ : From the $4(t+1)-1$ possible pairs $(s, s+1)$, there can be as many as $t+1$ such pairs that belong to the same defining set.
"Edge selection" refers to the process where we select edges $\left(v_{i_{1}}, v_{i_{2}}\right)$ that correspond to the popularity of an element in $S_{i_{1}}$ changed to that of an element in $S_{i_{2}}$. In order to ensure that the edges selected correspond to popularity changes of magnitude at most 1 , the following needs to hold: For any vertex, if an incoming directed arc into the vertex is selected, then an outgoing arc from the same vertex cannot be selected. Similarly, if for a vertex an outgoing arc is selected, no incoming arcs can be selected for that vertex.

Let $E^{\prime}$ stand for a set of arcs describing a collection of popularity changes of magnitude one. For $j \in\{1,2, \ldots, 2 t+2\}, S_{j}^{\prime}$ denotes the defining sets after the popularity changes induced by $E^{\prime}$. It is straightforward to see that because the trades are balanced, the change in discrepancy after the popularity updates equals $\sum_{j \in\{1,2, \ldots, t+1\}}\left|\sum_{i \in S_{2 j}^{\prime}} i-\sum_{i^{\prime} \in S_{2 j-1}^{\prime}} i^{\prime}\right|$, where

$$
\begin{equation*}
\sum_{j \in\{1,2, \ldots, t+1\}}\left|\sum_{i \in S_{2 j}^{\prime}} i-\sum_{i^{\prime} \in S_{2 j-1}^{\prime}} i^{\prime}\right| \leqslant 2\left|E^{\prime}\right| . \tag{1}
\end{equation*}
$$

We exhibit an edge set $E^{\prime}$ so that (1) holds with equality.
If $\left(v_{i_{1}}, v_{i_{2}}\right) \in E^{\prime}$, where $i_{1} \in\left\{2 j_{1}-1,2 j_{1}\right\}$ and $i_{2} \in\left\{2 j_{2}-\right.$ $\left.1,2 j_{2}\right\}$, then the following hold:

1) For any $v_{\ell} \in V,\left(v_{\ell}, v_{i_{1}}\right)$ and $\left(v_{i_{2}}, v_{\ell}\right)$ cannot both be included in $E^{\prime}$;
2) Let $S_{s}, s=i_{1} \pm 1$, be the companion defining set for $S_{i_{1}}$. For any $v_{\ell} \in V$, we cannot select $\left(v_{s}, v_{\ell}\right)$;
3) Let $S_{s}, s=i_{2} \pm 1$, be the companion defining set for $S_{i_{2}}$. For any $v_{\ell} \in V$, we cannot select $\left(v_{\ell}, v_{s}\right)$.
If we select $E^{\prime}$ according to (1)-(3), we will select at least 1 out of every 9 edges in $E$. Given that $|E| \geqslant 3\left(t+\frac{2}{3}\right)$, the resulting change in discrepancy is at least $2\left|E^{\prime}\right| \geqslant 2\left(\frac{|E|}{9}\right) \geqslant$ $\frac{6\left(t+\frac{2}{3}\right)}{9}$.
Due to space limitations, we omit a similar proof for the following sharper lower bound.

Theorem 10. Provided $19 \mid(t+1)$, the total set discrepancy is at least $\frac{14}{19}(t+1)$.

Lemma 11. There exists a minimal balanced trade with $v=16$ that has total set discrepancy $\leqslant 6$ for any collection of popularity swaps of magnitude $p=1$.

Proof: Choose the following defining sets, pictured in Figure 11 top:

$$
\begin{aligned}
& S_{1}=\{1,16\}, S_{2}=\{8,9\} ; S_{3}=\{4,5\}, S_{4}=\{2,7\} \\
& S_{5}=\{10,15\}, S_{6}=\{12,13\} ; S_{7}=\{3,14\}, S_{8}=\{6,11\}
\end{aligned}
$$

For any popularity swaps of $p=1$ involving elements in $S_{3}, S_{4}, S_{5}, S_{6}$, one has

$$
\begin{equation*}
\left|\sum_{j \in S_{i}^{\prime}} j-\sum_{j \in S_{i+1}^{\prime}} j\right| \leqslant 2 \tag{2}
\end{equation*}
$$

where $i \in\{3,5\}$. First consider $i=3$. Note that the adjacent swap $(4,5)$ does not change the balance. A simultaneous pair of adjacent swaps $(3,4)$ and $(4,5)$ is impossible, while the pair of adjacent swaps $(3,4)$ and $(5,6)$ leaves the sum unchanged. In this case, the only allowed adjacent swaps involving elements 2,7 would be $(1,2)$ and $(7,8)$, which when applied together leave the sum of elements in $S_{4}^{\prime}$ unchanged. The same is true for $i=5$. Hence, only one transposition per set is possible, leading to the upper bound above. Similarly,

$$
\begin{equation*}
\left|\sum_{j \in S_{1}^{\prime}} j-\sum_{j \in S_{2}^{\prime}} j\right| \leqslant 2 \tag{3}
\end{equation*}
$$

This follows since the only swaps that can affect 1 and 16 are $(1,2)$ and $(16,15)$, which in combination leave the sum unchanged. On the other hand, the sum of the second set, $S_{2}$, can change by at most 1 , either due to the transposition $(7,8)$ or the transposition $(9,10)$. As before, one can have at most one transposition per defining set, which produces the desired upper bound. A similar line of reasoning suggest that

$$
\begin{equation*}
\left|\sum_{j \in S_{7}^{\prime}} j-\sum_{j \in S_{8}^{\prime}} j\right| \leqslant 4 \tag{4}
\end{equation*}
$$

since it is possible to simultaneously increase/decrease both elements in both defining sets. Nevertheless, it is straightforward to show that if the above inequality is an equality, the second and third pair of defining sets do not change their sums, i.e., $S_{3}^{\prime}=\{4,6\}, S_{4}^{\prime}=\{3,7\}$ and $S_{5}^{\prime}=\{11,14\}, S_{6}^{\prime}=\{12,13\}$. A similar analysis can be used to examine the remaining choices for discrepancies of the first pair of defining sets. Hence, in the worst case, the discrepancy cannot exceed 6.
Given the large number of balanced defining sets, we only performed computer simulations for $v=12,16,20,24$ which show that the optimal total set discrepancies for these cases equal $6,6,8,10$, respectively. In particular, for $v=12$, the smallest possible change equals $2(t+1)$ as opposed to the value $3(t+1)$ described above. There are exactly 10 collections of balanced defining sets with discrepancy 6. Interestingly, for $v=16$ and $v=20$, the defining sets of minimum discrepancy under popularity changes of magnitude


Figure 1. Defining sets with smallest discrepancy for popularity changes of magnitude $p=1$ and $v=16$ (top) and $v=20$ (bottom).
$p=1$ are unique, and depicted in Figure 1 For $v=24$, there are exactly 22 optimal collections of defining sets.

To construct defining sets for for arbitrary values $v=$ $4(t+1) \geqslant 28$, one can "concatenate" the structures for $v=12,16,20,24$ by subdividing the elements of the sets into largest possible groupings of 24 or 20 or 16 elements and then grouping the remaining elements into one set. For example, the defining sets for $v=28$ can be constructed using the optimal patterns for defining sets with respect to elements in $\{1,2, \ldots, 16\}$ and $\{1,2, \ldots, 12\}$. In this case, one has to account for the discrepancy arising from swaps of elements at the "border" of the two groups (in this case, involving 16 and 17). It is easy to check that in this case, an additional discrepancy of 2 may arise, which for the given example gives rise to a total discrepancy of $6+6+2=14=2(t+1)$. Similarly, for $v=32$, we would group the first 20 and the last 12 elements together and obtain a total set discrepancy $8+6+2=16=2(t+1)$. Clearly, to reduce the boundary effects, one should maximize the size of groups of integers on which the individual patterns described above are to be used. A formal statement of this result follows.
Theorem 12. Let $v=4(t+1), m=\left\lfloor\frac{t+1}{6}\right\rfloor$ and $m^{\prime}=v-24 m$. Then there exists a balanced trade with block-size $k=2(t+1)$ and total set discrepancy

$$
12 m+\Delta=12\left\lfloor\frac{t+1}{6}\right\rfloor+\Delta
$$

where $\Delta=-2,2,4,6,6,8$ for $m^{\prime}=0,4,8,12,16,20$, respectively.
For $p=2$, we also have the following result.
Theorem 13. Let $v=4(t+1), m=\left\lfloor\frac{t+1}{5}\right\rfloor$ and $m^{\prime}=v-20 m$. Then there exists a balanced trade with block-size $k=2(t+1)$ and total set discrepancy

$$
26 m+\Delta=26\left\lfloor\frac{t+1}{5}\right\rfloor+\Delta
$$

where $\Delta=-4,4,8,12,18$ for $^{\prime}=0,4,8,12,16$, respectively. The result is a consequence of the fact that for $v=12,16,20$, the optimal total set discrepancies equal $12,18,22$, respectively. For $v=12$, there is unique defining sets of minimum discrepancy, while there are 12,7 optimal choices of defining sets for $v=16,20$, respectively.

Many open problems remain, pertaining to tighter bounds on discrepancies for $p \geqslant 1$, constructions of (optimal) defining sets of cardinality $\geqslant 2$ and others.

## REFERENCES

[1] Mehmet Fatih Aktaş, Amir Behruzi Far, Emina Soljanin, and Philip Whiting. Evaluating load balancing performance in distributed storage with redundancy. IEEE Transactions on Information Theory, 67(6):3623-3644, 2021.
[2] William M Brummond. Kirkman systems that attain the upper bound on the minimum block sum, for access balancing in distributed storage. arXiv preprint arXiv:1906.02157, 2019.
[3] Yeow Meng Chee, Charles J Colbourn, Hoang Dau, Ryan Gabrys, Alan CH Ling, Dylan Lusi, and Olgica Milenkovic. Access balancing in storage systems by labeling partial steiner systems. Designs, Codes and Cryptography, 88(11):2361-2376, 2020.
[4] L. Cherkasova and M. Gupta. Analysis of enterprise media server workloads: access patterns, locality, content evolution, and rates of change. IEEE/ACM Transactions on Networking, 12(5):781-794, 2004.
[5] Charles J Colbourn. Egalitarian steiner triple systems for data popularity. Designs, Codes and Cryptography, pages 1-23, 2021.
[6] Hoang Dau and Olgica Milenkovic. Maxminsum steiner systems for access balancing in distributed storage. SIAM Journal on Discrete Mathematics, 32(3):1644-1671, 2018.
[7] A. G. Dimakis, B. Godfrey, Y. Wu, M. J. Wainwright, and K. Ramchandran. Network coding for distributed storage systems. IEEE Transactions on Information Theory, 56(9):4539-4551, 2010.
[8] S. El-Rouayheb and K. Ramchandran. Fractional repetition codes for repair in distributed storage systems. In Proceedings of the 48th Annual Allerton Conference on Communication, Control, and Computing (Allerton), pages 1510-1517. IEEE, 2010.
[9] A Hedayat. The theory of trade-off for t-designs. In Coding Theory and Design Theory: Part II, Design Theory, page 101. 1990.
[10] HL Hwang. On the structure of (v, k, t) trades. Journal of Statistical Planning and Inference, 13:179-191, 1986.
[11] G. Joshi, Y. Liu, and E. Soljanin. On the delay-storage trade-off in content download from coded distributed storage systems. IEEE Journal on Selected Areas in Communications, 32(5):989-997, 2014.
[12] D. Leong, A. G. Dimakis, and T. Ho. Distributed storage allocations. IEEE Transactions on Information Theory, 58(7):4733-4752, 2012.
[13] Olgica Milenkovic and Kevin J Compton. Probabilistic transforms for combinatorial urn models. Combinatorics, Probability and Computing, 13(4-5):645-675, 2004.
[14] Ankit Singh Rawat, Dimitris S Papailiopoulos, Alexandros G Dimakis, and Sriram Vishwanath. Locality and availability in distributed storage. IEEE Transactions on Information Theory, 62(8):4481-4493, 2016.
[15] N. Silberstein and T. Etzion. Optimal fractional repetition codes based on graphs and designs. IEEE Transactions on Information Theory, 61(8):4164-4180, 2015.
[16] N. Silberstein and A. Gál. Optimal combinatorial batch codes based on block designs. Designs, Codes and Cryptography, 78(2):409-424, 2016.


[^0]:    This work was supported by the NSF grants CCF 1816913 (CJC) and 1814298 (OM). The authors acknowledge useful discussions with Dylan Lusi.

