

Permuting Data on Random-Access Block Storage*

Risi Thonangi
Duke University
rvt@cs.duke.edu

Jun Yang
Duke University
junyang@cs.duke.edu

ABSTRACT

Permutation is a fundamental operator for array data, with applications in, for example, changing matrix layouts and reorganizing data cubes. We consider the problem of permuting large quantities of data stored on secondary storage that supports fast random block accesses, such as solid state drives and distributed key-value stores. Faster random accesses open up interesting new opportunities for permutation. While external merge sort has often been used for permutation, it is an overkill that fails to exploit the property of permutation fully and carries unnecessary overhead in storing and comparing keys. We propose faster algorithms with lower memory requirements for a large, useful class of permutations. We also tackle practical challenges that traditional permutation algorithms have not dealt with, such as exploiting random block accesses more aggressively, considering the cost asymmetry between reads and writes, and handling arbitrary data dimension sizes (as opposed to perfect powers often assumed by previous work). As a result, our algorithms are faster and more broadly applicable.

1 Introduction

Permutation is a powerful primitive. In contrast to sorting, where the output address of a record depends on the results of comparing it with other records, here we are given a function that directly returns the output address of a record given its input address, without comparing its contents with other records. Below are some examples where we use permutation to reorganize data in useful ways.

Example 1. Suppose we have a large $p \times q$ dense matrix stored in a file in the row-major layout. To transpose it, or simply convert it into the column-major layout, we use a permutation that maps input address (i, j) to output address (j, i) . In terms of linearized addresses within the file (assuming all indices start from 0), this permutation moves the record at input address $i \times q + j$ to output address $j \times p + i$.

We can also convert the matrix into a blocked layout, where the matrix is stored as a grid of submatrices, say, in row-major order, and entries within each submatrix are stored, say, also in row-

major order. This layout is popular because it provides good locality of access for a wide range of matrix operations. For simplicity, suppose the matrix is $2^4 \times 2^2$. We can write an input address in the original row-major layout as a bit string $x_5x_4x_3x_2x_1x_0$, where $x_5x_4x_3x_2$ represents the row index in binary and x_1x_0 represents the column index in binary. A permutation mapping $x_5x_4x_3x_2x_1x_0$ to output address $x_5x_1x_4x_3x_2x_0$ converts the matrix into a blocked layout with a 2×2 grid of $2^3 \times 2$ submatrices.

Example 2. Suppose we have precomputed an aggregate for a sales data warehouse similar to TPC-H [11], where the group-by attributes include dimensions such as order-date, part-key, supplier-state, customer-state, etc., and the aggregated measure attributes include for example the total quantity of sales. The result may have been computed and stored in (order-date, part-key, supplier-state, customer-state, ...) order. Now we want to “resort” it in a different order (supplier-state, part-key, customer-state, order-date, ...), and use it to simultaneously compute other aggregates whose group-by attributes are prefixes of the new sort order. This type of reorganization is popular in computing multi-dimensional aggregates from data cubes (e.g., [1, 14, 8]).

We can regard this type of reorganization as a permutation of records in an address space formed by the dimensional attributes. Data can in fact be compactly stored as an array, as is commonly done in MOLAP [14, 13], with one-to-one mappings between integer array indices and dimensional attribute values. Records do not need to store either array indices or dimensional attribute values, as they are implied by the record addresses.

We study the problem of permuting lots of data (that cannot all fit in available memory) on secondary, block-based storage with fast random block accesses. Traditional hard disks are block-based, but random accesses are significantly slower than sequential accesses. On the other hand, modern flash-based storage devices can support random block accesses (reads and non-overwrite writes) as fast as sequential ones. Distributed key-value stores are another example designed for fast random accesses, where we can store a block as a key-value pair. Some of the hash-based key-value stores do not even provide a range scan interface. These storage solutions are quickly becoming popular alternatives to traditional hard drives for a variety of different reasons, ranging from fast speed, low energy consumption, to ease of maintenance (by being cloud-based).

For these storage solutions, it would seem natural to exploit their good support of random accesses in implementing permutation. But doing so effectively is still tricky because, as secondary storage, these storage solutions are still inevitably block-based. A simple, brute-force strategy would be to go through the output address space in order; for each output address, we would compute what input address should map to it, and simply read the record at that input address and write it out. However, this strategy generates lots

*This work is supported by the NSF award IIS-0916027 and an Innovation Research Program Award from HP Labs.

of small, random *record* reads, which are expensive on block-based storage where costs are charged by blocks.

Another strategy is sorting. We would associate each record with a key whose comparison order reflects the desired output order, and simply run the external merge sort algorithm on these key-carrying records. Sorting is strikingly effective. In fact, it has been shown [12] that *general* permutation is as expensive as sorting in the I/O model of computation (for large enough data). For many types of permutations encountered in practice, however, permutation can be cheaper than sorting when conditions are right.

We show, in this paper, that we can indeed do much better than sorting (and the brute-force strategy) for a large class of useful permutations. The key intuition is that, by analyzing the permutation function, we can figure out what subset of input records to read into memory at one time—we call them *action records*. For each set of action records, we permute them in memory and write them to desired output addresses. Each set of action records is chosen carefully such that they are clustered in both input and output address spaces; hence, reads and writes are on the granularity of blocks instead of records. To illustrate, consider a simple example.

Example 3. Consider a permutation that changes the ordering of records within each block and the ordering of blocks in the input file, but does not move any record across block boundaries. Suppose we have only one block of memory. We can choose records in each input block to be a set of action records. We read in each input block, permute its records, and write it out to the desired output block address. This strategy is “one-pass” in that it reads and writes each block only once (albeit in random order). Moreover, it remains one-pass regardless of how large the input is.

This feat would be difficult to accomplish with external merge sort, whose number of passes depends on the input size. With some tweaking, external merge sort can be made one-pass for special cases (e.g., almost-sorted input). However, its sequential access pattern makes changing block ordering hard, and it lacks a general framework for exploiting the properties of a given permutation.

Translating this intuition into a general, practical algorithm is not as simple. Some permutations require a multi-pass strategy to implement efficiently, with careful choices of action records for each pass. The realities of “imperfect” data dimension sizes and the performance characteristics of modern random-access block storage also make it challenging to achieve optimal performance in practice and for more general cases. Specifically, we make the following technical contributions:

- We handle *address-digit permutations* (ADP), which can be defined as functions permuting digits in a mixed-radix representation of the address space. ADP naturally captures all example permutations discussed earlier, as well as others such as bit-reversal (popular for Fast Fourier Transforms) and z-order (popular as a space-filling curve). We design efficient algorithms for performing ADP on random-access block storage, characterize the space of execution strategies, and develop techniques to search for optimal strategies that minimize cost.
- Many types of modern random-access block storage exhibit interesting asymmetries between read and write costs, and between non-overwrite and in-place write costs. For example, writes on flash-based storage tend to be more expensive than reads in terms of time, energy, as well as wear effect; in-place writes (and hence partial block writes) are even more expensive because they require first clearing a large “eraser block.” Writes on distributed key-value stores, and especially in-place writes, also tend to be more expensive than reads because of replication and consistency. Our approach is fully cognizant of

these performance characteristics. Our algorithms use “nice” I/O patterns that avoid partial block writes as much as possible. We build into our algorithms a mechanism called *filtered reads* to explore the trade-off between reads and writes—by performing more reads, we can reduce the number of passes and hence writes. Our optimization techniques consider this trade-off in finding execution strategies that minimize the overall cost.

- We show that our approach significantly outperforms external merge sort on several fronts. For typical block sizes, we only need a small amount of memory to enable efficient one-pass permutation of arbitrarily large input. If multiple passes are required, the number of passes depends on the inherent difficulty of the permutation, not on the input size; thus, the cost scales linearly with input size. In contrast, sorting will take more than one pass in most practical cases, and its cost is superlinear in the input size. Furthermore, sorting requires key comparisons and operates on a less efficient data representation, where each record must carry a key for comparison. For Example 2, the key includes all dimension attribute values. Our approach, however, does not require such a key to be stored, as the address of a record in the input already encodes where it would be in the output. This saving can be substantial with many dimensions, and it translates to proportionally smaller intermediate result sizes and lower cost per pass for our approach.
- While there has been some work on permuting data on secondary storage, our approach is practically more general and more efficient for random-access block storage. Previous approaches focus on hard drives, while ours achieves better performance by exploiting random block accesses more aggressively and considering the cost asymmetry between reads and writes. Furthermore, previous approaches typically assume that input, memory, and block sizes are all conveniently perfect powers of 2, and consider permutations defined using address bit permutations. For such approaches to work with an input address space consisting of d dimensions of arbitrary sizes to be permuted, padding is required and may blow up the data size (and hence execution cost) by a factor exponential in d . We choose to tackle general mixed-radix address spaces. While doing so requires addressing a number of nontrivial technical challenges, the resulting solutions are much more efficient and have much broader practical applicability.

After covering the preliminaries in Section 2, we will begin by considering a simple case in Section 3, where addresses can be represented using a single radix, and block and memory sizes are perfect powers of this radix. We will then consider the general, mixed-radix case in Section 4. In Sections 3 and 4, starting with a very simple algorithm, we will present a series of improvements and generalizations, each building on the previous. Section 5 presents our experimental evaluation. We discuss future extensions and related work in Section 6, before concluding in Section 7.

2 Preliminaries

We consider a class of permutations which we call *address-digit permutations* (ADP). It is related to the so-called *BPC permutations* [4], which we further discuss in Section 6.

Consider the (0-based) address of a record in a file in the mixed-radix representation with n digits. The n -dimensional space of such addresses is specified by integer radices $\tau = (r_{n-1}, \dots, r_1, r_0)$, where $r_i \geq 2$ for all $i \in [0, n)$. The lowest-order (least significant) digit is at position 0. We use $\langle x_{n-1}, \dots, x_1, x_0 \rangle_\tau$, where

$0 \leq x_i < r_i$, to denote the (linearized) record address

$$\sum_{i=0}^{n-1} \left(x_i \prod_{j<i} r_j \right) \quad (1)$$

within the file. We omit subscript τ when the context is clear.

An ADP is a permutation that can be defined as a permutation of digits in the mixed-radix representation of an address. More precisely, it is defined by a permutation π of the digit positions $0, 1, \dots, n-1$, which maps a digit at position i to position $\pi(i)$. This ADP moves the record at input address $\langle x_{n-1}, \dots, x_1, x_0 \rangle_\tau$ to output address

$\pi(\langle x_{n-1}, \dots, x_1, x_0 \rangle_\tau) = \langle x_{\pi^{-1}(n-1)}, \dots, x_{\pi^{-1}(1)}, x_{\pi^{-1}(0)} \rangle_{\pi(\tau)}$, where $\pi(\tau) = (r_{\pi^{-1}(n-1)}, \dots, r_{\pi^{-1}(0)})$ denotes the resulting radices for output record addresses.¹ As an example, consider π_1 in Figure 1 of a 6-digit input address space. It maps position 0 to 3, 1 to 1, etc. Effectively, π_1 moves a record from input address $\langle x_5, x_4, x_3, x_2, x_1, x_0 \rangle$ to output address $\langle x_4, x_3, x_0, x_2, x_1, x_5 \rangle$. Figure 1 also shows how the output is further permuted by another ADP π_2 . Although not shown here, the digits in general can have different radices, which permutations carry along.

Records are stored compactly in a file consisting of a sequence of *blocks*. We assume that all records have a fixed size, and no records span multiple blocks. Let B denote the block size, as measured by the number of records. We also assume that the data is *dense* in the input address space; i.e., there is a record at every address $\langle x_{n-1}, \dots, x_0 \rangle$, where $x_i \in [0, r_i)$ for each i . The total number of records is therefore $N = \prod_{i=0}^{n-1} r_i$. For example, Figure 2 illustrates an input file with a 6-bit address space, as well as the output file obtained by applying π_1 in Figure 1 (ignore for now the illustration of algorithmic steps between them).

Note that there is no need to store record addresses explicitly for permutation because they are implied by record positions within the file. For instance, as discussed, the data cube in Example 2 need not store any dimension attribute values for its records.

We briefly discuss how to handle variable-size records and non-dense input data in Section 6.

Additional Notation Let $S \subseteq [0, n)$ denote a set of digit positions in (the mixed-radix representation of) an address. A *setting* θ_S of S assigns each position $i \in S$ a value $\theta_S(i) \in [0, r_i)$. A setting θ_S is *complete* if $S = [0, n)$. Note that a specific record address is a complete setting. A setting θ_S is *partial* if $S \subsetneq [0, n)$. When writing a partial setting θ_S , we use \perp for digits whose positions are not in S . We write x^k for a sequence of k digits with the same value x . For example, $(1, 1, 0, 1, \perp^2)$ denotes a partial setting that sets the four highest-order digits (positions 5 through 2) to 1, 1, 0, and 1, while leaving the two lowest-order digits unset.

We say that two settings θ_{S_1} and θ_{S_2} are *compatible* if for every position $i \in S_1 \cap S_2$, $\theta_{S_1}(i) = \theta_{S_2}(i)$.

Given a setting θ_S , let $\Omega_{S'}(\theta_S)$, where $S' \subseteq S$, denote the setting for S' that is compatible with θ_S . Think of this operator as “project.” For example, $\Omega_{[2,4]}((1, 1, 0, 1, \perp^2)) = (\perp^2, 0, 1, \perp^2)$.

Let $\mathcal{U}_{S'}(\theta_S)$, where $S \subseteq S'$, denote the set of settings for S' that are compatible with θ_S . Think of this operator as “enumerate.” As a special case, $\mathcal{U}_{[0,n]}(\theta_S)$ denotes the set of all record addresses (i.e., complete settings) compatible with θ_S . As another special case, $\mathcal{U}_S(\langle \perp^n \rangle)$ denotes the set of all possible settings for S ; we use \mathcal{U}_S as a shorthand for it. For example, if all radices are 2, $\mathcal{U}_{[0,6]}$ returns all bit strings of length 6. Finally, we let $\mathcal{U}_\emptyset = \{\langle \perp^n \rangle\}$.

Given an ADP π and a set S of digit positions in the input address space, we use $\pi(S)$ to denote the set $\{\pi(i) \mid i \in S\}$, the resulting

¹We overload π so that in addition to as a function acting on digit positions, we can also interpret it as a function on addresses or a function on radices.

positions of the digits in S in the output address space. Given a setting θ_S , we use $\pi(\theta_S)$ to denote the result of applying π to θ_S , which is a setting of $\pi(S)$ in the output record address that assigns each digit position $j \in \pi(S)$ the value $\theta_S(\pi^{-1}(j))$.

3 Single-Radix ADP with Perfect-Power Block and Memory Sizes

Let us start simple. Assume that all radices in an n -digit input address space are the same (r), so the total number of input records is $N = r^n$. Assume further that the block size is $B = r^b$, a perfect power of the radix; and that the size of available memory (also measured in number of records) is $M = r^m$, where $b \leq m \leq n$.

We need some new terminology to help with exposition. We call the b lowest-order digits (at positions $b-1, \dots, 1, 0$) *in-digits* (for in-block digits); they specify the address of a record within the block storing it. We call the remaining higher-order digits (at positions $n-1, n-2, \dots, b$) *out-digits* (to contrast with in-digits). A setting of all out-digits identifies a block within the file; we call this setting the *block id*. For example, in Figure 2, $r = 2$ and $n = 6$, so addresses are 6-bit strings; $b = 3$, so a block can be identified by the $n - b = 3$ highest-order bits (out-digits) while the remaining $b = 3$ bits (in-digits) vary among records within the block; $m = 4$, so the memory holds $2^m = 16$ records or $2^{m-b} = 2$ blocks.

Given an ADP π , we are especially interested in those out-digits in the input address space that are turned by π into in-digits in the output address space. We call these the *entering* digits of π . Let $F(\pi) = \{i \in [b, n) \mid \pi(i) < b\}$ denote their digit positions in the input address space (“ F ” is for “entering F rom positions”). Let $\delta(\pi) = |F(\pi)|$ denote the number of entering digits. For example, for π_1 in Figure 1, there is only one entering digit 5 (position within the input address space). Thus, $\delta(\pi_1) = 1$. On the other hand, if we consider the composition of two ADPs in Figure 1, $\pi_2 \circ \pi_1$ (i.e., π_1 followed by π_2), then $F(\pi_2 \circ \pi_1) = \{4, 5\}$, and $\delta(\pi_2 \circ \pi_1) = 2$.

As we shall see, given an ADP π , $\delta(\pi)$ gives a measure of how “difficult” π is. Example 3 has discussed the case when $\delta(\pi) = 0$; i.e., π keeps all out-digits out and in-digits in (though π may change the ordering among out-digits and the ordering among in-digits). In this case, we can do π in one pass with only one block of memory.

3.1 Basic Single-Radix ADP

Basic One-Pass Single-Radix ADP With more than one block of memory, we can handle a more difficult ADP π with $\delta(\pi) > 0$ in one pass. More precisely, we show how, using memory of size r^m (or, equivalently, r^{m-b} blocks), we can perform in one pass an ADP π where the number of entering digits $\delta(\pi) \leq m - b$.

On a high level, $\text{ADP}_{\text{pass}}^{\text{basic}}$ (shown as Algorithm 1) works in iterations. Each iteration processes a set of “action records”: it first reads $r^{\delta(\pi)}$ blocks containing these action records into memory (they fit because $\delta(\pi) \leq m - b$), permutes these records in memory, and writes them out as $r^{\delta(\pi)}$ output blocks.

More specifically, $\text{ADP}_{\text{pass}}^{\text{basic}}$ selects and processes action records for each iteration as follows. It partitions the input digit positions into two groups:

- *Action* digits, A , which include in-block digit positions $[0, b)$ and entering digit positions $F(\pi)$. For example, Figure 1 shows the actions digits for π_1 and π_2 .
- *Non-action* digits, \bar{A} , which include all remaining digits.

Each iteration of $\text{ADP}_{\text{pass}}^{\text{basic}}$ considers a particular setting $\theta_{\bar{A}}$ of \bar{A} (Line 4). Figure 2, which shows how $\text{ADP}_{\text{pass}}^{\text{basic}}$ works for π_1 in Figure 1, serves as an example accompanying the description below:




Figure 1: Illustration of ADPs π_1 , π_2 , and their composition $\pi_2 \circ \pi_1$, in a 6-digit address space. Action digits chosen by $\text{ADP}_{\text{pass}}^{\text{basic}}$ for π_1 and π_2 (assuming $b = 3$) are shown—action digit positions in the input address space are marked with stripes to their right; in the output address space they are marked to their left.




Figure 2: $\text{ADP}_{\text{pass}}^{\text{basic}}$ performing π_1 in Figure 1; $r = 2$, $n = 6$, $b = 3$, and $m = 4$. Records are color-coded by their output ordering. Block boundaries are shown with thick lines. Action digits in input and output addresses are shown in red. There are a total of 4 iterations. Contents of the memory during these iterations are shown with the two middle rows; dashed lines separate the iterations.

Algorithm 1: $\text{ADP}_{\text{pass}}^{\text{basic}}(R, \pi)$: basic single-pass, single-radix ADP.

Input: R : a file of r^n records; π : an ADP where $\delta(\pi) \leq m - b$.
Output: a file containing permuted records.

- 1 $A \leftarrow [0, b) \cup F(\pi)$;
- 2 $\bar{A} \leftarrow [0, n) \setminus A$;
- 3 $\pi_{\text{mem}} \leftarrow g_2 \circ \pi \circ g_1$, where:
 $g_1 \leftarrow i \mapsto (i + 1)$ -th smallest element of A , for $0 \leq i < |A|$, and
 $g_2 \leftarrow i \mapsto$ number of elements in $\pi(A)$ smaller than i ;
- 4 **foreach** $\theta_{\bar{A}} \in \mathcal{U}_{\bar{A}}$ **do** // each possible setting of \bar{A}
- 5 clear memory;
- 6 **foreach** $\theta_{[b, n)} \in \mathcal{U}_{[b, n)}(\theta_{\bar{A}})$ in asc. order **do**
 | read input block with id $\theta_{[b, n)}$ and append it to memory;
- 7 permute records in memory by π_{mem} ;
- 8 **foreach** $\vartheta_{[b, n)} \in \mathcal{U}_{[b, n)}(\pi(\theta_{\bar{A}}))$ in asc. order **do**
 | write the next memory block as output block with id $\vartheta_{[b, n)}$;

- The action records are those whose input addresses are compatible with $\theta_{\bar{A}}$. There are exactly $r^{|\bar{A}|}$ of them, and they occupy exactly $r^{|\bar{A}| - b} = r^{\delta(\pi)}$ full input blocks. $\text{ADP}_{\text{pass}}^{\text{basic}}$ reads all these blocks into memory (Lines 6–7).

For example, in Figure 2, where $\bar{A} = \{3, 4\}$, the first iteration reads two blocks (memory size) of action records whose input addresses are compatible with $\langle \perp, 0, 0, \perp^3 \rangle$.

- Initially, the ordering of action records in memory is consistent with their ordering in the input file. $\text{ADP}_{\text{pass}}^{\text{basic}}$ permutes them in memory, in place, such that their ordering becomes consistent with the desired output ordering. This in-memory permutation, π_{mem} , is itself an ADP over the memory address space. It is invariant across iterations, and is precomputed by Line 3.

Figure 2 illustrates how this π_{mem} is applied to action records in memory consistently across iterations.

- Now, we go through memory blocks in order, and write them to appropriate locations in the output file (Lines 9–10). The output blocks have ids that compatible with the setting for \bar{A} (or more precisely, $\pi(\theta_{\bar{A}})$) in the current iteration of the outer loop.

For example, in the first iteration shown in Figure 2, $\theta_{\bar{A}} = \langle \perp, 0, 0, \perp^3 \rangle$, so $\pi(\theta_{\bar{A}}) = \langle 0, 0, \perp, \perp^3 \rangle$, which gives us the two output blocks whose ids’ two highest-order bits are 0s.

There are a total of $r^{|\bar{A}|}$ iterations. For example, in Figure 2, it takes four iterations to apply π_1 .

Overall, $\text{ADP}_{\text{pass}}^{\text{basic}}$ reads the entire input once in a specific order, and writes the output in a specific order. Thus, the I/O cost of $\text{ADP}_{\text{pass}}^{\text{basic}}$, if we simply count the total number of block accesses, is exactly $2N/B$. Intuitively, $\text{ADP}_{\text{pass}}^{\text{basic}}$ permutes digits within A (which includes in-digits in either input or output address space) in memory in each of its iterations; it permutes out-digits simply using random block I/Os.

Basic Multi-Pass Single-Radix ADP $\text{ADP}_{\text{pass}}^{\text{basic}}$ allows up to $m - b$ out-digits in the input address to move to in-digit positions in the output address. Given an ADP π , if $\delta(\pi) > m - b$, we cannot complete π using a single invocation of $\text{ADP}_{\text{pass}}^{\text{basic}}$. In this case, we adopt a simple multi-pass strategy: we use $\lceil \frac{\delta(\pi)}{m-b} \rceil$ passes, each of which calls $\text{ADP}_{\text{pass}}^{\text{basic}}$ to move up to $m - b$ out-digits to in-digit positions. As a shorthand, we say that a pass *admits* an out-digit in the input address if the pass moves this digit to an in-digit position in the output address.

For example, consider the ADP $\pi = \pi_2 \circ \pi_1$ in Figure 1. Suppose $m = 4$ and $b = 3$. Since $m - b < \delta(\pi) = 2$ (both 4 and 5 are entering input digit positions), we cannot perform π in one pass of $\text{ADP}_{\text{pass}}^{\text{basic}}$. Instead, we would use two passes, as illustrated by Figure 1: the first pass performs π_1 (to admit x_5), and the second pass perform π_2 (to admit x_4).

Discussion The basic approach outlined in this subsection already offers compelling advantages over sorting. The cost of multi-pass ADP, measured by the total number of block I/Os, is $\lceil \frac{\delta(\pi)}{m-b} \rceil \cdot \frac{2N}{B}$. The number of passes, $\lceil \frac{\delta(\pi)}{m-b} \rceil$, is independent of the input size. Note that $\delta(\pi) \leq b$ because the number of entering digits cannot exceed the number of in-digits. Therefore, as long as $m \geq 2b$, we are guaranteed a one-pass solution. For example, assuming a typical block size of 8kB and a fairly small record size of 4B, we need only $(\frac{8\text{KB}}{4\text{B}})^2 \cdot 4\text{B} = 16\text{MB}$ of memory to guarantee a one-pass solution for arbitrarily large input and arbitrary ADP. Bigger records would lower the memory requirement further. In contrast, one-pass sorting is generally impossible once the input exceeds memory size, and the cost of multi-pass sorting is superlinear in the input size.

These advantages are possible because of $\text{ADP}_{\text{pass}}^{\text{basic}}$ ’s aggressive use of random block I/Os. To make our approach more practical, however, we still need to consider additional performance characteristics of modern random-access block storage, as well as arbitrary radices that do not “fit” memory and block sizes; both issues can significantly impact efficiency and applicability.

3.2 Improved Single-Radix ADP

We consider characteristics of modern random-access block storage to improve the basic approach of Section 3.1 in several ways. First, we can make I/Os “nicer”—instead of accessing random individual blocks, we strive to make larger requests spanning multiple consecutive blocks, and turn writes into appends. While they do not change the number of block I/Os in the case we consider in this section (single-radix ADP with perfect-power block and memory sizes), nicer I/Os can be more efficient if per-request overhead is high. Furthermore, they will become crucial for bounding the number of block I/Os in the general case we consider in Section 4, where I/O requests are no longer aligned with block boundaries.

The second idea, already mentioned in Section 1, is the use of “filtered reads” to account for the cost asymmetry between reads and writes. Intuitively, at the expense of reading some unnecessary data, filtered reads let us admit more digits in each pass, thereby decreasing the required number of passes (and hence writes).

In this subsection we show how to realize these improvements. We first present, in Section 3.2.1, an improved one-pass algorithm, which follows a given “plan” specifying what and how to permute. The plan choice affects performance (e.g., admitting more digits beyond $m - b$ will incur higher read costs). Thus, in Section 3.2.2, we discuss how to choose the optimal strategy for performing a given ADP, which may involve multiple passes and choosing a plan for each pass to minimize overall cost.

3.2.1 Improved One-Pass Single-Radix ADP

$\text{ADP}_{\text{pass}}^{\text{sr}}$ (shown as Algorithm 2) is the improved version of $\text{ADP}_{\text{pass}}^{\text{basic}}$. We describe the improvements in terms of the ideas outlined at the beginning of this subsection; the corresponding modifications to $\text{ADP}_{\text{pass}}^{\text{basic}}$ are color-coded in Algorithm 2.

Before starting, we note a high-level change: $\text{ADP}_{\text{pass}}^{\text{sr}}$ takes two extra input arguments (A, β) , which define the plan to be followed in execution. A is the set of action digits, while β specifies the size of input segments, which we will explain shortly. We will discuss how to choose (A, β) intelligently in Section 3.2.2.

Making Reads Bigger $\text{ADP}_{\text{pass}}^{\text{basic}}$ reads individual blocks. With careful choices of action records, however, we may be able to make fewer, bigger read requests in the unit of *input segments*. In general, an input segment contains all records that share a common input address prefix (i.e., higher-order digit values) and, therefore, lie consecutively in the input file. Plan parameter $\beta \geq b$ specifies the size of each input segment to be r^β . The r^β records in a segment lie in $r^{\beta-b}$ consecutive input blocks. **Line 9** reads each input segment.

Making Writes Bigger We can do the same for writes. Instead of writing individual blocks, we can write one *output flush* at a time. Analogous to an input segment, an output flush contains all output records that share the same output address prefix and, therefore, lie consecutively in the output file. Given an ADP π and the set of action digits A , we determine the output flush size as follows:

$$\phi(\pi, A) = \max\{j \mid [0, j] \subseteq \pi(A)\}. \quad (2)$$

In other words, we look for the longest output address suffix (i.e., lowest-order digits) consisting entirely of action digits. Let $\varphi = \phi(\pi, A)$ denote the thus computed *flush (digit) position* (**Line 1**). The size of each output flush is r^φ , and all r^φ in this flush lie in $r^{\varphi-b}$ consecutive output blocks. **Line 14** writes each output flush.

For example, consider π_1 in Figure 1. Examining the destination positions of action digits in the output of π_1 , we see that the longest suffix containing only action digits has length 4, so $\varphi = 4$. Thus, we can write $2^{\varphi-b} = 2$ blocks (which happen to be all of memory) as a single output flush. Figure 2 clearly illustrates this possibility.

Why do we compute φ by Eq. (2)? Given π and A , we want output flushes to be as large as possible. Recall that all action records in memory at the same time share the same setting for non-action digits, but a flush contains records whose output addresses have different settings for digit positions in $[0, \varphi)$. Therefore, the range $[0, \varphi)$ cannot contain any non-action digits; Eq. (2) computes the largest possible φ under this constraint.

Turning Writes into Appends In addition to making writes bigger, we can in fact ensure that the output file can be written as a number of *partitions*, such that we only need to *append* flushes to these partitions. To reduce the number of partitions that need to be kept open simultaneously for appends, we produce partitions in

partition groups. With each memory full of action records, we append a flush to each of the partitions in the same partition group. After completing a group of partitions in this round-robin fashion, we move on to the next group. Bigger partitions and fewer partitions per group are nicer.

Like a flush, a partition consists of all records sharing a common output address prefix. Given π and A , we determine the output partition size as follows:

$$\rho(\pi, A) = \max\{j \leq n \mid [\phi(\pi, A), j] \cap \pi(A) = \emptyset\}. \quad (3)$$

Intuitively, in the output address space, starting from the flush position φ and heading toward higher-order positions, we skip over any non-action digits and stop at the first action digit we encounter (or n , if no action digit exists in $[\varphi, n)$). Let $\varrho = \rho(\pi, A)$ denote the thus computed *partition (digit) position* (**Line 1**). The size of each partition is r^ϱ , and there are a total of $r^{n-\varrho}$ partitions. We will justify Eq. (3) shortly.

For example, consider π_1 in Figure 1. Examining the destination positions of action digits in the output of π_1 , we see that no action digit exists in $[\varphi, n) = [4, 6)$, so $\varrho = 6$, meaning that the entire output file is one partition. Next, consider π_2 . In the output address space of π_2 , we see $\varphi = 3$. We skip position 3 as it is not an action digit, and stop at position 4, which is. Therefore, $\varrho = 4$, meaning that there are $2^{n-\varrho} = 4$ partitions of size $2^\varrho = 16$ each.

Note that we can consider a setting for the output digit positions $[\varrho, n)$ as a *partition number*, which is a 0-based sequence number that identifies a partition within the output file. Each partition contains $r^{\varrho-\varphi}$ flushes. We can consider a setting for the output digit positions $[\varphi, \varrho)$ as a *flush number*, which is a 0-based sequence number that identifies a flush within a partition.

To implement group-at-a-time, append-only flushing, we need to pick an appropriate order for bringing in sets of action records. The outermost loop of $\text{ADP}_{\text{pass}}^{\text{basic}}$ simply steps through all possible settings for \bar{A} , leaving the ordering unspecified. In $\text{ADP}_{\text{pass}}^{\text{sr}}$, we break this loop down into a nest of two.² Note that we can divide non-action digits (more precisely, their positions in the output address space) into two sets: non-action partition number digits, i.e., $\pi(\bar{A}) \cap [\varrho, n)$; and the remaining non-action digits, which are exactly the flush number digits $[\varphi, \varrho)$ (since no digit in \bar{A} enters $[0, \varphi)$ by the definition of φ).

- The outer loop of $\text{ADP}_{\text{pass}}^{\text{sr}}$ (**Line 4**) iterates over the partition groups by enumerating every possible setting $\vartheta_{\pi(\bar{A}) \cap [\varrho, n)}$ for non-action partition number digits. Each group thus consists of partitions whose numbers are compatible with this setting. This definition is natural, because in the outer loops, we do not have settings for action digits.
- Then, for each partition group, the inner loop (**Line 5**) enumerates all settings for \bar{A} compatible with $\vartheta_{\pi(\bar{A}) \cap [\varrho, n)}$ in order, which basically involves stepping through flush numbers and extending $\vartheta_{\pi(\bar{A}) \cap [\varrho, n)}$ with them to form settings for \bar{A} .

In every iteration of the inner loop on **Line 5**, the output loop (**Line 13**) writes $r^{|\pi(A) \setminus [0, \varphi)|}$ flushes. Each flush is effectively appended to its respective partition in the current partition group being processed. It is straightforward to confirm that the number of partitions in each group is also $r^{|\pi(A) \setminus [0, \varphi)|}$.

Having seen how the algorithm operates, we can now justify our calculation of partition size in Eq. (3). Consider the output address space. Flush number digits, i.e., $[\varphi, \varrho)$, should not come

²Actually, an appropriately ordered loop would suffice, and there is no need for this algorithm to compute ϱ , but the presentation is more clear with two loops.

Algorithm 2: $\text{ADP}_{\text{pass}}^{\text{sr}}(R, \pi, A, \beta)$: improved single-pass, single-radix ADP.

Input: R : a file of r^n records; π : an ADP (where $\delta(\pi) > m - b$ is possible); $A \subseteq [0, n)$: input digit positions identifying the action digits to be used by the algorithm; $\beta \in (0, n)$: an input digit position specifying the input segment size to be used by the algorithm.

Output: a file containing permuted records.

- 1 $\varphi \leftarrow \phi(\pi, A)$; $\varrho \leftarrow \rho(\pi, A)$; // using Eq. (2) and (3)
- 2 $\bar{A} \leftarrow [0, n) \setminus A$;
- 3 $\pi_{\text{mem}} \leftarrow g_2 \circ \pi \circ g_1$, where:
 $g_1 \leftarrow i \mapsto (i + 1)$ -th smallest element of A , for $0 \leq i < |A|$, and
 $g_2 \leftarrow i \mapsto$ number of elements in $\pi(A)$ smaller than i ;
- 4 **foreach** $\vartheta_{\pi(\bar{A}) \cap [0, n)} \in \mathcal{U}_{\pi(\bar{A}) \cap [0, n)}$ **do** // each partition group
 - 5 **foreach** $\vartheta_{\pi(\bar{A})} \in \mathcal{U}_{\pi(\bar{A})}(\vartheta_{\pi(\bar{A}) \cap [0, n)})$ in asc. order **do** // each flush #
 - 6 $\theta_{\bar{A}} \leftarrow \pi^{-1}(\vartheta_{\pi(\bar{A})})$;
 - 7 clear memory;
 - 8 **foreach** $\theta_{[\beta, n)} \in \mathcal{U}_{[\beta, n)}(\Omega_{\bar{A} \cap [\beta, n)}(\theta_{\bar{A}}))$ in asc. order **do**
 - 9 $E \leftarrow$ read the segment of input records with addresses compatible with $\theta_{[\beta, n)}$;
 - 10 **foreach** $e \in E$ with input address compatible with $\theta_{\bar{A}}$, in order, **do**
 - 11 \lfloor append e to memory;
 - 12 permute records in memory by π_{mem} ;
 - 13 **foreach** $\vartheta_{[\varphi, n)} \in \mathcal{U}_{[\varphi, n)}(\vartheta_{\pi(\bar{A})})$ in asc. order **do**
 - 14 \lfloor write, as one flush, the next r^φ records in memory to output addresses compatible with $\vartheta_{[\varphi, n)}$;

from action digits, because we need to step through flush numbers on Line 5 but cannot control settings for actions digits. Under this constraint, Eq. (3) computes the largest possible ϱ , in order to maximize the partition size.

Permuting More Using Filtered Reads Recall that $\text{ADP}_{\text{pass}}^{\text{basic}}$ can admit only $m - b$ digits, so an ADP π with $\delta(\pi) > m - b$ cannot be performed in one pass. We now show how to make $\text{ADP}_{\text{pass}}^{\text{sr}}$ perform π in one pass. In this pass, $\text{ADP}_{\text{pass}}^{\text{sr}}$ writes the output exactly once, as before. However, $\text{ADP}_{\text{pass}}^{\text{sr}}$ can use “filtered reads”—when it reads from an input segment, it may keep only a filtered subset of the records in memory as action records. The net effect is more reads, but greater flexibility that enables us to permute more.

A plan (A, β) with filtered reads is one where the range $[0, \beta)$ of input address digit positions contain some non-action digits; i.e., $[0, \beta) \cap A \neq \emptyset$. Let $A^+ = [0, \beta) \setminus A \subseteq \bar{A}$ denote the set of non-action digits in $[0, \beta)$. In $\text{ADP}_{\text{pass}}^{\text{basic}}$ we had $[0, b) \subseteq A$, which can be considered as a special case where $A^+ = \emptyset$.

We just need some small changes to enable filtered reads. For each possible setting $\theta_{\bar{A}}$ of \bar{A} , our goal is still to read in all records with compatible addresses as action records. In the loop starting on Line 8, we still read input by segments. Note that we need to project $\theta_{\bar{A}}$ by $\Omega_{\bar{A} \cap [\beta, n)}(\cdot)$ (Line 8), because if $A^+ \neq \emptyset$, $\bar{A} \not\subseteq [\beta, n)$.

More importantly, whenever we read an input segment, we get records with all possible settings for $[0, \beta)$. But we only want a subset as action records—namely, those with addresses compatible with $\theta_{\bar{A}}$, or, more precisely, addresses with the same setting as $\theta_{\bar{A}}$ for digits in A^+ . Therefore, we only keep one in every $r^{|A^+|}$ records; others are discarded. We thus call $r^{|A^+|}$ the *waste factor*. This filtering step is reflected by the loop on Line 10 (which we actually implement together with Line 9 in a streaming fashion).

Note that for each $\theta_{\bar{A}}$, $\text{ADP}_{\text{pass}}^{\text{sr}}$ executes Line 9 $r^{|A| - \beta + |A^+|}$ times, because there are exactly $|A^+|$ digits in $[0, \beta)$ that are not in A . Thus, $\text{ADP}_{\text{pass}}^{\text{sr}}$ performs a total of $r^{|A|} r^{|A| - \beta + |A^+|} = r^{n - \beta + |A^+|}$ segment reads. The total number of input segments is $r^{n - \beta}$, meaning that input is read $r^{|A^+|}$ times—consistent with the waste factor.

As a concrete example, Figure 3 shows how, using filtered reads, $\text{ADP}_{\text{pass}}^{\text{sr}}$ is able to handle $\pi_2 \circ \pi_1$ from Figure 1 in one pass, which was not possible for $\text{ADP}_{\text{pass}}^{\text{basic}}$. Here, $A^+ = \{2\}$, so the waste

Algorithm 3: $\text{Optimize}_{\text{pass}}^{\text{sr}}(\pi)$: find the best one-pass plan for a single-radix ADP, which may involve filtered reads.

Input: π : an ADP.
Output: a plan (A, β) with the lowest cost.

- 1 $A \leftarrow [0, b) \cup F(\pi)$; $\beta \leftarrow b$;
- 2 **while** $|A| < m$ **do** // make bigger output flushes with all available memory
 - 3 $\lfloor A \leftarrow A \cup \{\pi^{-1}(\min([0, n) \setminus \pi(A)))\}$;
- 4 **while** $|A| > m$ **do** // ensure action records fit in memory
 - 5 $\lfloor A \leftarrow A \setminus \{\pi^{-1}(\rho(\pi, A)) \text{ if } \rho(\pi, A) < n \text{ else } \max(\pi(A))\}$;
- 6 $\beta \leftarrow \max\{i \in [b, n) \mid [b, i) \subseteq A\}$; // make bigger input segments
- 7 **return** (A, β) ;

factor is 2. Therefore, we end up reading each input block twice (and throwing away half of the data). However, doing so nets us a one-pass strategy that reads all records twice and writes them once, which is strictly better than a two-pass strategy that would read them twice and write them twice.

As special case, if $A^+ = \emptyset$, all input records pass filtering; the waste factor becomes 1; and $\text{ADP}_{\text{pass}}^{\text{sr}}$ reads each input segment exactly once.

Requirement and Cost Analysis $\text{ADP}_{\text{pass}}^{\text{sr}}$ assumes that its given plan (A, β) meets the following requirements: $|A| \leq m$ (i.e., action records fit in memory); $\beta \geq b$ (i.e., input segments are no smaller than blocks); $\phi(\pi, A) \geq b$ (i.e., output flushes are no smaller than blocks). In Section 3.2.2, we ensure that we only generate plans meeting these requirements.

Because we assume both block and memory sizes to be perfect powers of the single radix, all input segments and output flushes align with block boundaries. Thus, we use a simple but effective cost model counting block I/Os. The number of blocks written by $\text{ADP}_{\text{pass}}^{\text{sr}}$ is $\frac{N}{B}$. The number of blocks read is $\frac{N}{B} r^{|A^+|}$, where $r^{|A^+|}$ is the waste factor. To account for the asymmetry between read and write costs, we model the total cost as a linear combination of numbers of blocks read and written, where parameter $\alpha \geq 0$ represents the cost of reads relative to writes. Recalling that $A^+ = [0, \beta) \setminus A$, the cost of $\text{ADP}_{\text{pass}}^{\text{sr}}$ given plan (A, β) is

$$\text{cost}(A, \beta) = \frac{N}{B} \left(\alpha 2^{\lfloor [0, \beta) \setminus A \rfloor} + 1 \right). \quad (4)$$

3.2.2 Improved Single-Radix ADP with Optimization

Given an ADP π as well as the amount of memory $M = r^m$, we need to find a strategy for performing π with the lowest cost. Although we can always compete π in one pass using $\text{ADP}_{\text{pass}}^{\text{sr}}$ with filtered reads, the best strategy may involve multiple passes instead. Before showing how to find such overall best strategies, however, we discuss how to choose the best single-pass plan.

Choosing Plan for One-Pass Single-Radix ADP Given π , we choose the best one-pass plan (A, β) using $\text{Optimize}_{\text{pass}}^{\text{sr}}$ (Algorithm 3). It starts with a “baseline” plan with $A = [0, b) \cup F(\pi)$ and $\beta = b$, the same choice as $\text{ADP}_{\text{pass}}^{\text{basic}}$. There are two cases.

The first case is when $\delta(\pi) \leq m - b$. Here, the baseline plan already meets the memory and segment/flush size requirements without using filtered reads. It is arguably optimal for the cost model in Eq. (4). However, if $\delta(\pi) < m - b$, there is memory to accommodate more digits in A . We choose to add those digits that immediately follow $[0, \phi(\pi, A))$ in the output address space (Line 3). Doing so makes flushes bigger (even though this improvement is not modeled by Eq. (4)).

The second case is when $\delta(\pi) > m - b$. The baseline choice of A exceeds the memory requirement, so we must remove $\delta(\pi) - (m - b)$ digits from A . But which ones? We cannot remove any digit in $F(\pi)$ from A —that would introduce a non-action digit in positions $[0, b)$ in the output address space, making the output flush smaller

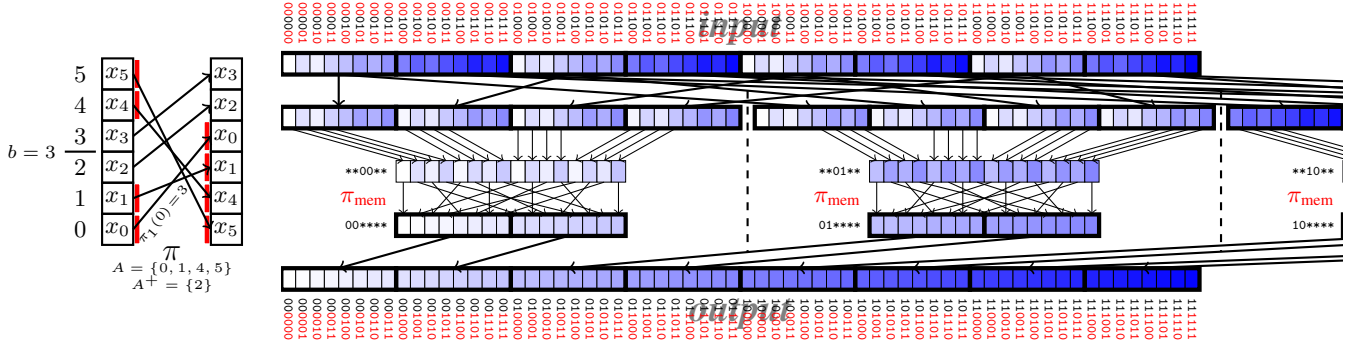


Figure 3: $\text{ADP}_{\text{pass}}^{\text{sr}}$ performing $\pi = \pi_2 \circ \pi_1$ from Figure 1 in one pass using filtered reads; $r = 2$, $n = 6$, $b = 3$, and $m = 4$ (same as Figure 2). There are a total of 4 iterations separated by dashed lines; the last two are clipped. The second row shows the input blocks read by each iteration; keep in mind that they are simply streamed in and filtered to produce the memory contents shown on the third row.

than a block. Therefore, we can only remove from A input address digit positions in $[0, b)$. The removed digits become A^+ , so $|A^+| = \delta(\pi) - m + b$. Which digits we remove from A does not affect the cost computed by Eq. (4). However, we prefer removing the action digit at the current partition position $\rho(\pi, A)$ in the output address space (Line 5), which makes partitions bigger by increasing $\rho(\pi, A)$. If $|A^+|$ is still too big when the entire input is one partition (i.e., $\rho(\pi, A) = n$), we remove the higher-order action digits in the output address space, keeping $\rho(\pi, A) = n$.

Finally, we increase β as much as possible without affecting A or A^+ (Line 6). In other words, we make input segments bigger without affecting the cost in Eq. (4).

The analysis above leads us directly to the following lemma:

Lemma 1. *Give an ADP and the available memory size, $\text{Optimize}_{\text{pass}}^{\text{sr}}$ finds an optimal one-pass plan for $\text{ADP}_{\text{pass}}^{\text{sr}}$ (possibly with filtered reads) under the cost model of Eq. (4). For an ADP π with $\delta(\pi) = \delta$, the cost of this optimal plan is*

$$\text{cost}(\delta) = (n - b) \left(\alpha 2^{\min\{0, \delta - (m - b)\}} + 1 \right).$$

Overall Optimization of Single-Radix ADP Given π , $\text{ADP}_{\text{pass}}^{\text{sr}}$ (Algorithm 4) finds the optimal (possibly multi-pass) strategy and executes it by invoking $\text{ADP}_{\text{pass}}^{\text{sr}}$ with an appropriate plan for each pass. To optimize, we first establish an upper bound on the optimal number of passes to search through. Then, for each particular number of passes, we consider only the strategy that balances the numbers of digits admitted in each pass as much as possible.

For example, suppose $b = 16$, $m = 19$, and $\delta(\pi) = 14$. Without filtered reads, each pass can admit $m - b = 3$ digits. We will need at most 5 passes. Among the four-pass strategies, $\text{ADP}_{\text{pass}}^{\text{sr}}$ will consider the one that admits (4, 4, 3, 3) digits in each of the four passes. Less balanced strategies, e.g., (3, 3, 3, 5) and (4, 4, 4, 2), will not be considered.

The following theorem establishes that this optimization method ensures the optimality of $\text{ADP}_{\text{pass}}^{\text{sr}}$ under reasonable assumptions about the cost function (see appendix for proof). Indeed, the function $\text{cost}(\delta)$ in Lemma 1 meets these assumptions.

Theorem 1. *Suppose $\text{cost}(\delta)$ is the cost to perform an ADP π with $\delta(\pi) = \delta$ in one pass. If $\text{cost}(\delta)$ is convex in δ and flat for $\delta \in [1, m - b]$, then $\text{ADP}_{\text{pass}}^{\text{sr}}$ finds the optimal number of passes and the optimal number of digits to admit in each of the passes.*

Strategies considered by $\text{ADP}_{\text{pass}}^{\text{sr}}$ are strict improvement and superset of the those considered by our basic approach in Section 3.1. Therefore, we inherit all advantages of the basic approach over sorting discussed at the end of Section 3.1: having a number of passes independent of input size and a running time linear in it, as well

Algorithm 4: $\text{ADP}_{\text{pass}}^{\text{sr}}(R, \pi)$: improved multi-pass, single-radix ADP.

Input: R : a file of r^n records; π : an ADP.
Output: a file containing permuted records.

```

1  $\Delta_{\text{opt}} \leftarrow \perp$ ;  $C_{\text{opt}} \leftarrow \infty$ ;
2 for  $q = \lceil \frac{\delta(\pi)}{m-b} \rceil$  to 1 step  $-1$  do // search for the optimal number of passes
   // compute the list  $\Delta$  of the number of out-digits to move
   // to in-digit positions in each of the  $q$  passes:
3    $\Delta \leftarrow \emptyset$ ;  $C \leftarrow 0$ ;  $q_1 \leftarrow \delta(\pi) \bmod q$ ;
4   for  $p = 1$  to  $q_1$  do
5      $\delta_p \leftarrow (\delta(\pi) - q_1) / q + 1$ ;  $\Delta.append(\delta_p)$ ;  $C \leftarrow C + \text{cost}(\delta_p)$ ;
6   for  $p = q_1 + 1$  to  $q$  do
7      $\delta_p \leftarrow (\delta(\pi) - q_1) / q$ ;  $\Delta.append(\delta_p)$ ;  $C \leftarrow C + \text{cost}(\delta_p)$ ;
8   if  $C \leq C_{\text{opt}}$  then
9      $\Delta_{\text{opt}} \leftarrow \Delta$ ;  $C_{\text{opt}} \leftarrow C$ ;
10 for  $p = 1$  to  $|\Delta_{\text{opt}}|$  do // in each pass, move  $\delta_p$  out-digits to in-digit positions
11    $\delta_p \leftarrow$  the  $p$ -th element of  $\Delta_{\text{opt}}$ ;
12    $D_{\text{enter}} \leftarrow \pi^{-1}$  (the  $\delta_p$  smallest elements of  $\pi(F(\pi))$ );
13    $D_{\text{exit}} \leftarrow \pi^{-1}$  (the  $\delta_p$  smallest elements of  $T(\pi)$ );
14    $L \leftarrow [0, b) \cup D_{\text{enter}} \setminus D_{\text{exit}}$  ordered by  $i \mapsto \pi(i)$ ;
15    $L.append([b, n) \setminus D_{\text{exit}} \cup D_{\text{enter}}$  ordered by  $i \mapsto \pi(i)$ );
16    $\pi_p \leftarrow i \mapsto$  (0-based) position of  $i$  in  $L$ , for  $0 \leq i < n$ ;
17    $(A, \beta) \leftarrow \text{Optimize}_{\text{pass}}^{\text{sr}}(\pi_p)$ ;
18    $R \leftarrow \text{ADP}_{\text{pass}}^{\text{sr}}(R, \pi_p, A, \beta)$ ;
19    $\pi \leftarrow \pi_p^{-1} \circ \pi$ ; // compute the "remaining" permutation
20 return  $R$ ;
```

as being able to do one-pass permutation with very small memory requirement. $\text{ADP}_{\text{pass}}^{\text{sr}}$ also goes further: it has nicer I/O patterns, and it can exploit the cost asymmetry between reads and writes with filtered reads. Interestingly, if this asymmetry is so large (i.e., $\alpha \rightarrow 0$) that reads are practically free, $\text{ADP}_{\text{pass}}^{\text{sr}}$ will choose to complete any ADP in one pass, using essentially an improved version of the brute-force strategy discussed in Section 1.

4 General ADP

We now consider the general case where the input address space has mixed radices $\tau = (r_{n-1}, \dots, r_0)$, and block size B and memory size M may not be perfect powers of a single radix. Recall from Section 2 that we assume data are dense in the original address space, so the number of input records is $N = \prod_{i=0}^{n-1} r_i$. However, as we will soon see, it is often necessary to “decompose” the original address space in order to perform a permutation, and data may become no longer dense in an intermediate address space—i.e., some seemingly valid addresses may not have records associated with them. One of the challenges we tackle in this section is how to keep records compactly stored in files without physically padding these conceptual holes in the address space.

We first present, in Section 4.1, a general one-pass algorithm,

which builds on $\text{ADP}_{\text{pass}}^{\text{sr}}$. We then motivate the need for address space decomposition and present our solution in Section 4.2. Finally, in Section 4.3, we discuss how to optimize a general ADP, which may require a multi-pass strategy.

4.1 One-Pass ADP

ADP_{pass} (shown as Algorithm 5) is our general one-pass ADP algorithm. It uses the same ideas and follows the same flow as $\text{ADP}_{\text{pass}}^{\text{sr}}$ (Section 3.2.1). We process one memory full of action records at a time: they are read in from input segments, filtered if needed, permuted in memory, and then written out as output flushes. We go through the sets of action records in a way such that output is produced in groups of partitions: for each group, we append flushes to the partitions in this group in a round-robin fashion. For additional explanation and examples, see Section 3.2.1. Here, we focus on the differences (highlighted in Algorithm 5) from $\text{ADP}_{\text{pass}}^{\text{sr}}$.

First, ADP_{pass} takes two additional input arguments, \mathfrak{T} and ϖ , beyond π , A , and β . (\mathfrak{T}, ϖ) specifies the input address space, which in general can be a decomposed one. $(\mathfrak{T}, \varpi \circ \pi^{-1})$ specifies the output address space after applying ADP π . We will explain these notations further in Section 4.2; for now, it suffices to remember what they present respectively.

To deal with the generality of mixed-radix, decomposed address spaces, we also need a few helper functions:

- $\text{maxSize}^{\mathfrak{T}, \varpi}(S)$ returns the maximum number of records whose input addresses are compatible with a setting of digits not in S .
- $\text{outSize}^{\mathfrak{T}, \varpi \circ \pi^{-1}}(\vartheta_{[p, n]})$, where $\vartheta_{[p, n]}$ is a prefix setting of output address, returns the exact number of records whose output addresses are compatible with $\vartheta_{[p, n]}$.
- $\text{offset}_P^{\mathfrak{T}, \varpi \circ \pi^{-1}}(\vartheta)$, where ϑ is an output address and P is a prefix of output digits, returns the number of records whose output addresses precede ϑ and are compatible with $\Omega_P(\vartheta)$ (i.e., they share the same prefix as ϑ).

Computing these functions is simple assuming dense data—e.g., $\text{maxSize}^{\mathfrak{T}, \varpi}(S)$ would be the product of the radices of the digits in S . In the general case, however, calculation becomes non-trivial; we will show how to do it in Section 4.2, when we discuss decomposed address spaces in depth.

In $\text{ADP}_{\text{pass}}^{\text{sr}}$ (and $\text{ADP}_{\text{pass}}^{\text{basic}}$), we specified how to permute a set of action records in memory using an ADP (π_{mem}) over the memory address space. In ADP_{pass} , we choose a more explicit implementation (which is simpler to describe for general, decomposed address spaces). Suppose we need to process a set of action records, whose addresses share the same setting $\vartheta_{\pi(\bar{A})}$ for non-action digits. In preparation, we divide available memory into chunks (Lines 8–9), one for each flush that will be appended to one of the partitions in the current group. The exact sizes of these flushes (and corresponding memory chunks) are calculated using the helper function $\text{outSize}(\cdot)$. When reading input, for each action record we encounter, we use its output address to compute which output flush and what offset within this flush it should go to, and simply write it to that offset within the corresponding memory chunk (Line 14). The offset is calculated by the helper function $\text{offset}(\cdot)$. After we read in all action records, they are already in the desired order within each memory chunk; we simply write these memory chunks out as output flushes (Line 16).

A more fundamental difference between ADP_{pass} and the less general $\text{ADP}_{\text{pass}}^{\text{sr}}$ —but largely hidden from the algorithm pseudo code—is that input segments and output flushes may no longer align with block boundaries. Since the underlying storage is block-based and partial block writes are especially expensive, we take

Algorithm 5: $\text{ADP}_{\text{pass}}(R, \mathfrak{T}, \varpi, \pi, A, \beta)$: general single-pass ADP.

Input: (\mathfrak{T}, ϖ) : an n -digit permuted decomposed input address space; R : a file of records ordered according to (\mathfrak{T}, ϖ) ; π : an ADP; $A \subseteq [0, n)$: input digit positions identifying the action digits to be used by the algorithm; $\beta \in (0, n)$: an input digit position specifying the input segment size to be used by the algorithm.

Output: a file containing permuted records, ordered according to $(\mathfrak{T}, \varpi \circ \pi^{-1})$.

- 1 $\varphi \leftarrow \phi(\pi, A)$; $\varrho \leftarrow \rho(\pi, A)$; // using Eq. (2) and (3)
- 2 $\bar{A} \leftarrow [0, n) \setminus A$;
- 3 **reserve** $\text{maxSize}^{\mathfrak{T}, \varpi}(\pi^{-1}([0, n) \cap A))$ memory blocks as output buffer;
- 4 **foreach** $\vartheta_{\pi(\bar{A}) \cap [0, n)} \in \mathcal{U}_{\pi(\bar{A}) \cap [0, n)}$ **in asc. order do** // each partition group
 - 5 **foreach** $\vartheta_{\pi(\bar{A})} \in \mathcal{U}_{\pi(\bar{A})}(\vartheta_{\pi(\bar{A}) \cap [0, n)})$ **in asc. order do** // each flush #
 - 6 $\theta_{\bar{A}} \leftarrow \pi^{-1}(\vartheta_{\pi(\bar{A})})$;
 - 7 **clear** memory allocated for chunks;
 - 8 **foreach** $\vartheta_{[\varphi, n)} \in \mathcal{U}_{[\varphi, n)}(\vartheta_{\pi(\bar{A})})$ **in asc. order do**
 - 9 **allocate** a chunk of memory of size $\text{outSize}^{\mathfrak{T}, \varpi \circ \pi^{-1}}(\vartheta_{[\varphi, n)})$ that corresponds to output addresses compatible with $\vartheta_{[\varphi, n)}$;
 - 10 **foreach** $\theta_{[\beta, n)} \in \mathcal{U}_{[\beta, n)}(\Omega_{\bar{A} \cap [\beta, n)}(\theta_{\bar{A}}))$ **in asc. order do**
 - 11 $E \leftarrow$ read the segment of input records with addresses compatible with $\theta_{[\beta, n)}$;
 - 12 **foreach** $e \in E$ with input address compatible with $\theta_{\bar{A}}$, **in order, do**
 - 13 $\vartheta \leftarrow \pi(\text{input address of } e)$;
 - 14 **store** e at (0-based) offset $\text{offset}_{[\varphi, n)}^{\mathfrak{T}, \varpi \circ \pi^{-1}}(\vartheta)$ within the memory chunk corresponding to $\Omega_{[\varphi, n)}(\vartheta)$;
 - 15 **foreach** $\vartheta_{[\varphi, n)} \in \mathcal{U}_{[\varphi, n)}(\vartheta_{\pi(\bar{A})})$ **in asc. order do**
 - 16 **write**, as one flush, the chunk in memory corresponding to $\vartheta_{[\varphi, n)}$ to output addresses compatible with $\vartheta_{[\varphi, n)}$;

some additional steps to ensure practical efficiency. To avoid partial block writes, we reserve a block of memory as an output buffer for each partition in a partition group (Line 3). The number of buffer blocks (i.e., number of partitions in a group) is calculated using the helper function $\text{maxSize}(\cdot)$. Each output flush first goes through the buffer block for the corresponding partition; an actual append happens whenever the buffer block fills up (or the partition is complete). For simplicity, Algorithm 5 does not show all details. These buffer blocks avoid partial block writes that would have been incurred by flushes that do not begin or end on block boundaries. These buffer blocks also allow us to use flushes that are smaller than blocks, which is more flexible than $\text{ADP}_{\text{pass}}^{\text{sr}}$. Of course, memory required by the buffer blocks must be accounted for and balanced against memory used for action records, which we consider next and in Section 4.3.

Requirement and Cost Analysis Suppose ADP_{pass} is invoked with a plan $(\mathfrak{T}, \varpi, \pi, A, \beta)$. Its memory consumption can be calculated below, as the sum of the number of action records (the first term) and the total size of buffer blocks (the second term):

$$\text{mem}(\mathfrak{T}, \varpi, \pi, A, \beta) = \text{maxSize}^{\mathfrak{T}, \varpi}(A) + B \cdot \text{maxSize}^{\mathfrak{T}, \varpi}(\pi^{-1}([\rho(\pi, A), n)) \cap A). \quad (5)$$

ADP_{pass} assumes that its plan meets the following requirements:

- $\text{mem}(\mathfrak{T}, \varpi, \pi, A, \beta) \leq M$ (memory requirement);
- $\text{maxSize}^{\mathfrak{T}, \varpi}([0, \beta)) \geq B$ (i.e., input segments are no smaller than blocks);
- $\text{maxSize}^{\mathfrak{T}, \varpi \circ \pi^{-1}}([0, \rho(\pi, A))) \geq B$ (i.e., output partitions are no smaller than blocks—a weaker requirement than $\text{ADP}_{\text{pass}}^{\text{sr}}$, thanks to output buffer blocks).

In Section 4.3, we ensure that we only generate plans meeting these requirements.

As with $\text{ADP}_{\text{pass}}^{\text{sr}}$, we model the cost of ADP_{pass} as follows, where C_r is the number of block reads, C_w is the number of block writes, and $\alpha \geq 0$ represents the cost of reads relative to writes:

$$\text{cost}(\overline{x}, \overline{w}, \pi, A, \beta) = C_r + \alpha C_w. \quad (6)$$

However, calculating C_r and C_w is significantly more involved than the less general case of $\text{ADP}_{\text{pass}}^{\text{sr}}$.

First, $C_r = (\frac{N}{B} + U_r) \times \text{maxSize}^{\overline{x}, \overline{w}}([0, \beta] \setminus A)$. Here, U_r denotes the number of input segment boundaries that do not coincide with block boundaries.³ The multiplier $\text{maxSize}^{\overline{x}, \overline{w}}([0, \beta] \setminus A)$ is the waste factor of filtered reads (1 if reads are not filtered). In the general case, where input segment boundaries generally do not coincide with block boundaries, we use the number of input segments, $\text{maxSize}^{\overline{x}, \overline{w}}([0, \beta], n)$, for U_r .

Next, $C_w = \frac{N}{B} + U_w$. Here U_w denotes the number of output partition boundaries that do not coincide with block boundaries. In the general case, where input segment boundaries generally do not coincide with block boundaries, we use the number of output partitions, $\text{maxSize}^{\overline{x}, \overline{w}}(\pi^{-1}([\pi(A), n]))$, for U_w . Recall that unaligned flush boundaries do not incur additional block writes because of output buffer blocks.

4.2 Address Decomposition

With arbitrary radices and block and memory sizes, there are now situations where we cannot pick a valid set of action digits A to perform an ADP (even with filtered reads). For example, consider transposing a big $p \times q$ matrix, where $p > M$ and $q > M$. If we model the address space using two digits (row and column numbers), including any digit in A at all would exceed the memory requirement of ADP_{pass} , because every radix is bigger than M .

Conceptually, the solution is straightforward—we simply *decompose* a digit into multiple “smaller” components such that the product of the radices of these components equals the radix of the original digit. Then, the task of permuting bigger digits can be accomplished by permuting smaller components. In the case of matrix transpose, for example, if p and q are perfect powers of 2, we can decompose them into $\log_2 p$ and $\log_2 q$ bits, respectively.

In practice, however, we might not be able to find such perfect decompositions. For example, p and q may not be perfect powers; worse, both can be prime numbers. In such cases, we have to deal with *imperfect* decompositions, where the original radix does not equal the product of component radices. Imperfect decompositions complicate address manipulations because data may not be dense in the decomposed address spaces; Eq. (1) no longer works. For example, consider imperfect decompositions in Figure 4. We see that there are (shaded) “gaps” in the decomposed address spaces. For instance, address $\langle 0, 1, 2 \rangle$ in $T^{3,(2,3)}$ has no record associated with it, even though all digit values are less than their respective radices; however, the same address $\langle 0, 1, 2 \rangle$ is valid in $T^{(3,2),3}$, a decomposition with the same radices as $T^{3,(2,3)}$ but nonetheless different. Eq. (1) happens to work for $T^{(3,2),3}$, but not for $T^{3,(2,3)}$.

One way to cope with imperfect decompositions is padding—we simply fill gaps in the decomposed address space with dummy records, essentially blowing up the original address space. For example, a 1025×513 matrix could be padded and handled as a 2048×1024 matrix if we want a decomposed address space consisting of bits. However, dummy records take space in files, making each pass more expensive; injecting dummy records into the original input and removing them from the final output also incur extra cost. Worse, the space overhead of padding is exponential in the

³Intuitively, segment and block boundaries divide the input file into parts. Each part incurs a block read, and there are a total of $\frac{N}{B} + U_r$ parts.

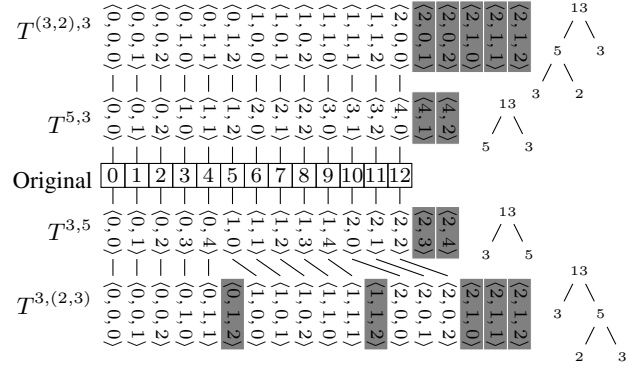


Figure 4: Different (imperfect) decompositions of the same single-digit address space with radix 13. The superscripts of T indicate radices of the decomposed address spaces. $T^{3,(2,3)}$ is obtained by further decomposing the lower-order digit of $T^{3,5}$; $T^{(3,2),3}$ is obtained by further decomposing the higher-order digit of $T^{5,3}$. Invalid addresses are shaded.

number of digits—given an n -digit original address space, if all digits need to be decomposed and if, for each digit, valid addresses make up for a fraction $f < 1$ of the address space decomposed from this digit, then only f^n fraction of the addresses in the overall decomposed address space are valid. Therefore, we need a better approach for making ADP practical in the general case.

In this subsection, we present an approach that keeps records stored compactly at all times, even if the address space is decomposed imperfectly and data become no longer dense. The notion of address decomposition requires careful formulation—as Figure 4 shows, two decomposed address spaces ($T^{(3,2),3}$ and $T^{3,(2,3)}$) can come from the same original space, have the same list of radices, but require very different address mappings. In the following, we first study what it means to decompose a single digit, and then discuss how to decompose and permute address spaces in general.

4.2.1 Decomposing a Single Digit

A *decomposition of a digit* with radix r into two digits with radices $r'' > 1$ and $r' > 1$ is *valid* if $r''r' \geq r$, $(r'' - 1)r' < r$, and $r''(r' - 1) < r$. Under this decomposition, value x (radix r) is mapped to a pair of values x'' (radix r'') and x' (radix r'), where $x'' = \lfloor x/r' \rfloor$ and $x' = x \bmod r'$.

A *recursive decomposition* of an *original digit* with radix r into a sequence of *final digits* is a node-labeled binary tree T , where:

- Each node $u \in T$ represents a digit, and is labeled by $\text{rad}(u)$, its radix. The root of T , denoted $\text{root}(T)$, represents the original digit. The sequence of leaves in T , denoted $\text{Leaves}(T)$, represents the sequence of final digits, from the highest order to the lowest.
- A node $u \in T$ having left child u'' and right child u' represents the decomposition of the digit corresponding to u into those corresponding to u'' and u' . We call u'' (u') the *higher-order* (*lower-order*, resp.) *child* of u .

Indeed, Figure 4 shows the decompositions as trees, whose structures allow us to distinguish them.

We begin with a crucial observation characterizing the valid settings in a decomposed address space.

Lemma 2. Consider an η -digit address space with radices $(r_{\eta-1}, \dots, r_0)$ obtained from a single digit with radix r through decomposition T . Let $\theta_{[p,\eta]}$ denote a valid setting for digits at positions $[p, \eta]$ ($1 \leq p \leq \eta$). The set X of valid addresses that are compatible with $\theta_{[p,\eta]}$ can be partitioned into subsets according to their values of the digit at position $p-1$. Let $X_\lambda \subseteq X$ denote the subset

of addresses with value λ for digit position $p - 1$. Let Z_0 denote the set of valid addresses in T compatible with $\langle 0^{\eta-p+1}, \perp^{p-1} \rangle$. Then:

1. The set of possible values for λ is $[0, t)$ for some $t \in (0, r_{p-1}]$; i.e., there are no “holes.”
2. X_0, X_1, \dots, X_{t-2} are identical to Z_0 modulo their values for digit positions $[p - 1, \eta)$.
3. X_{t-1} is a subset of Z_0 modulo digit positions $[p - 1, \eta)$; i.e., for every address in X_{t-1} , replacing its values for digit positions $[p - 1, \eta)$ by 0’s yields an address in Z_0 .

Lemma 2 reveals a small but important difference from the special case where the entire decomposition T is perfect (i.e., every binary decomposition in T is perfect). With a perfect T , the set of valid values for digit position $p - 1$ is always $[0, r_{p-1})$; however, for general T , the set of valid values at digit position $p - 1$ depends on the prefix setting $\theta_{[p, \eta)}$. For example, for $T^{3, (2, 3)}$ in Figure 4, given prefix $\langle 0, 0, \perp \rangle$, the lowest-order digit can be 0, 1, or 2; however, given prefix $\langle 1, 1, \perp \rangle$, it can only be 0 or 1.

Lemma 2 prompts us to introduce a useful quantity when working with decomposed address spaces. Given a prefix setting $\theta_{[p, \eta)} = \langle x_{\eta-1}, x_{\eta-2}, \dots, x_p, \perp^p \rangle$ for digit positions $[p, \eta)$, let $\kappa^T(\theta_{[p, \eta)})$ denote the number of valid settings for $[0, \eta)$ compatible with $\theta_{[p, \eta)}$. For example, for $T^{3, (2, 3)}$ in Figure 4, $\kappa^{T^{3, (2, 3)}}(\langle 0, 0, \perp \rangle) = 3$ and $\kappa^{T^{3, (2, 3)}}(\langle 1, \perp, \perp \rangle) = 5$. For $T^{(3, 2), 3}$, $\kappa^{T^{(3, 2), 3}}(\langle 1, \perp, \perp \rangle) = 6$.

We can calculate $\kappa^T(\cdot)$ in time linear in the number of digits:

Lemma 3. *Suppose T is a recursive decomposition of a single digit into an η -digit address space. Given $\theta_{[p, \eta)}$, $\kappa^T(\theta_{[p, \eta)})$ can be computed in $O(|T|)$.*

More specifically, we calculate $\kappa^T(\cdot)$ as follows. First, we can compute $\kappa^T(\cdot)$ for all-0 prefix settings in top-down fashion on T :

Lemma 4. *Suppose T is a recursive decomposition of a single digit into an η -digit address space. Define $z(u)$ for each node $u \in T$ as follows:*

- If u is $\text{root}(T)$, $z(u) = 1$;
- If u has children u'' (higher-order) and u' (lower-order), then $z(u'') = z(u) \cdot \text{rad}(u')$ and $z(u') = z(u)$.

Then, for any leaf $u \in \text{Leaves}(T)$, $z(u) = \kappa^T(\langle 0^{\eta-i}, \perp^i \rangle)$, where i is the position of the digit that u corresponds to in the final decomposed address space.

Next, we show how to calculate $\kappa^T(\cdot)$ for an arbitrary prefix setting iteratively, starting from its setting for the highest-order digit and expanding to lower-order digits one at a time. This calculation uses $\kappa^T(\cdot)$ for all-0 prefix settings (calculated by Lemma 4).

Lemma 5. *Suppose T is a recursive decomposition of a single digit with radix r into an η -digit address space. Then,*

- $\kappa^T(\langle \perp^\eta \rangle) = r$.
- $\kappa^T(\langle x_{\eta-1}, \dots, x_{i+1}, x_i, \perp^i \rangle) = \begin{cases} z_i & \text{if } x_i < \lceil \frac{s_{i+1}}{z_i} \rceil - 1 \text{ or } s_{i+1} \equiv 0 \pmod{z_i} \\ s_{i+1} \bmod z_i & \text{otherwise,} \end{cases}$

where $s_{i+1} = \kappa^T(\langle x_{\eta-1}, \dots, x_{i+1}, \perp^{i+1} \rangle)$, and $z_i = \kappa^T(\langle 0^{\eta-i}, \perp^i \rangle)$.

4.2.2 Decomposing and Permuting an Address Space

A decomposition \mathfrak{T} of an m -digit original address space is a sequence of node-labeled binary trees $(T_{m-1}, \dots, T_1, T_0)$, where each T_i ($i \in [0, m)$) is a recursive decomposition of the digit at position i in the original input address. The original digit for T_i has radix $\text{rad}(\text{root}(T_i))$. We denote the number of trees in \mathfrak{T} by $|\mathfrak{T}|$; i.e., $|\mathfrak{T}| = m$. We denote the number of digits after decomposition by $\|\mathfrak{T}\|$; i.e., $\|\mathfrak{T}\| = \sum_{i=0}^{|\mathfrak{T}|} |\text{Leaves}(T_i)|$.

A *permuted decomposed address space* is defined by a pair (\mathfrak{T}, ϖ) , where \mathfrak{T} specifies the decomposition from the original address space and ϖ maps a digit position $i \in [0, \|\mathfrak{T}\|)$ to a leaf (of one of the trees) in \mathfrak{T} . Thus, the radices of (\mathfrak{T}, ϖ) are $(\text{rad}(\varpi(\|\mathfrak{T}\| - 1)), \dots, \text{rad}(\varpi(0)))$. As notational shorthands, let $\varpi^{\text{tree}}(i)$ denote the tree $T \in \mathfrak{T}$ where $\varpi(i) \in T$, and let $\varpi^{\text{pos}}(i)$ denote number of leaves following $\varpi(i)$ in $\varpi^{\text{tree}}(i)$.

We say that a permuted decomposed address space (\mathfrak{T}, ϖ) is *consistent* if ϖ preserves the ordering of leaves within each tree in \mathfrak{T} ; i.e., for all $k > k'$ where $\varpi^{\text{tree}}(k) = \varpi^{\text{tree}}(k')$, $\varpi^{\text{pos}}(k) > \varpi^{\text{pos}}(k')$. We consider only consistent permuted decomposed address spaces (our optimization procedure in Section 4.3 ensures that it is the case). In the following, unless otherwise noted, all permuted decomposed address spaces are consistent.

Consider a permuted decomposed address space (\mathfrak{T}, ϖ) . We call a subset P of digit positions $[0, \|\mathfrak{T}\|)$ a *component-wise prefix* (or *suffix*) if for every $T \in \mathfrak{T}$, those digits in P that correspond to the leaves of T , if any, form a prefix (or suffix, resp.) of $\text{Leaves}(T)$.⁴ It is easy to see that every prefix (or suffix) of the digit positions in the permuted decomposed address space is also a component-wise prefix (or suffix, resp.), but not vice versa.

Suppose θ_P is a component-wise prefix (or suffix) setting for (\mathfrak{T}, ϖ) . Let $\Omega_T(\theta_P)$, where $T \in \mathfrak{T}$, denote a setting for the address space consisting of only $\text{Leaves}(T)$ as digits, obtained by removing from θ_P all digit i where $\varpi^{\text{tree}}(i) \neq T$. Clearly, $\Omega_T(\theta_P)$ will be a prefix (or suffix, resp.) setting.

We now extend the definition of $\kappa(\cdot)$ to the permuted decomposed address space (\mathfrak{T}, ϖ) . Suppose θ_P is a component-wise prefix setting for this address space. Let $\kappa^{\mathfrak{T}, \varpi}(\theta_P)$ denote the number of valid addresses compatible with θ_P . This quantity can be readily computed in a component-wise fashion, as the product of number of possible suffix settings in the decomposition of each of the original digits (recall that these numbers can be computed by Lemma 3).

Theorem 2. *Suppose (\mathfrak{T}, ϖ) is a consistent decomposed address space and θ_P is a component-wise prefix setting in this space. Then $\kappa^{\mathfrak{T}, \varpi}(\theta_P) = \prod_{T \in \mathfrak{T}} \kappa^T(\Omega_T(\theta_P))$.*

We now revisit ADP_{pass} in Section 4.1 and fill in the missing details. The input arguments (\mathfrak{T}, ϖ) together specify the permuted decomposed input address space, which may be a decomposition of the original input address or an intermediate one (if ADP_{pass} is called by a multi-pass strategy). Note that after applying π to this address space, the output address space will be specified by $(\mathfrak{T}, \varpi \circ \pi^{-1})$. The helper functions can be computed as follows.

Lemma 6. *Consider an n -digit input address space (\mathfrak{T}, ϖ) , and an output address space (\mathfrak{T}, ϖ') obtained by applying ADP π to the input address space, where $\varpi' = \varpi \circ \pi^{-1}$.*

- Let S be a component-wise suffix of the input address digits. Then $\text{maxSize}^{\mathfrak{T}, \varpi'}(S) = \kappa^{\mathfrak{T}, \varpi}(\Omega_{[0, n) \setminus S}(\langle 0^n \rangle))$.

⁴More precisely, for every $T \in \mathfrak{T}$, there exists no $u, v \in \text{Leaves}(T)$ where u precedes (or follows, resp.) v in $\text{Leaves}(T)$, but $\varpi^{-1}(u) \notin P$ and $\varpi^{-1}(v) \in P$.

- Let $\vartheta_{[p,n]}$ be a prefix setting of the output address. Then, $\text{outSize}^{\mathfrak{S}, \varpi'}(\vartheta_{[p,n]}) = \kappa^{\mathfrak{S}, \varpi'}(\vartheta_{[p,n]})$.
- Let $\langle x_{n-1}, \dots, x_0 \rangle$ be an output address, and $[p, n]$ be a prefix of output digits. Then, $\text{offset}_{[p,n]}^{\mathfrak{S}, \varpi'}(\langle x_{n-1}, \dots, x_0 \rangle) = \sum_{i=0}^{p-1} x_i s_i$, where $s_i = \kappa^{\mathfrak{S}, \varpi'}(\langle x_{n-1}, \dots, x_{i+1}, 0, \perp^i \rangle)$.

It remains to be shown how to implement the operator $\mathcal{U}(\cdot)$ (used throughout ADP_{pass}) efficiently in a permuted decomposed input address space. That is, given a component-wise prefix setting θ_Q , we need to enumerate all valid settings for a “bigger” component-wise prefix P (where $P \supseteq Q$) in ascending order. For this enumeration, it can be shown that we can compute, for each digit, when to carry over to the next digit, as follows:

Lemma 7. Suppose P and $P \cup \{k\}$ ($k \notin P$) are component-wise prefixes in an n -digit address space (\mathfrak{S}, ϖ) . Given a setting θ_P , the maximum value for digit k among valid addresses compatible with θ_P is

$$\left\lceil \frac{\kappa^T(\Omega_T(\theta_P))}{\kappa^T(\Omega_T(\langle 0^{n-k}, \perp^k \rangle))} \right\rceil, \quad (7)$$

where $T = \varpi^{\text{tree}}(k)$.

However, it would be inefficient to pay $O(|T|)$ time to calculate $\kappa(\cdot)$ for each digit and for every setting enumerated. Instead, we offer an efficient implementation of $\mathcal{U}(\cdot)$ that incrementally maintains relevant $\kappa(\cdot)$ results during enumeration. The algorithm Enumerate (shown as Algorithm 6) is expressed as a continuation, using a Python-style **yield/return** syntax. It precomputes and remembers in \mathbf{z} all $\kappa(\cdot)$ results ever used in the denominator of Eq. (7), as they do not change during enumeration. For $\kappa(\cdot)$ results used in the numerator of Eq. (7)—there is one for each digit in P given the current setting—Enumerate remembers and update them incrementally across invocations.

Enumerate has an amortized running time of $O(1)$ per setting enumerated, because updating each digit takes constant time, and the average number of digit updated is $O(1)$ per setting enumerated (the lowest-order digit changes once per setting, but higher-order digit changes exponentially less frequently).

4.3 General ADP with Optimization

Given input data in a mixed-radix address space, we now discuss how to carry out a given ADP, with block size B and memory size M . We assume the cost model in Section 4.1. In general, an optimal strategy may require multiple passes: each pass may choose to decompose its input address space, and may employ filtered reads. The original address space may undergo multiple decompositions and permutations.⁵

Before presenting the algorithms in this section, we introduce a procedure $\text{Decomp}(\mathfrak{S}, \varpi, S)$ for enumerating decompositions of (\mathfrak{S}, ϖ) (plus itself). Input argument S specifies a subset of the digits in (\mathfrak{S}, ϖ) to be potentially decomposed. A decomposition returned by $\text{Decomp}(\mathfrak{S}, \varpi, S)$ either leaves each digit S as is, or decomposes it into two digits. $\text{Decomp}(\mathfrak{S}, \varpi, S)$ enumerates all such possible decompositions. We use \star to denote a decomposition returned by $\text{Decomp}(\mathfrak{S}, \varpi, S)$. Let $(\mathfrak{S}_\star, \varpi_\star)$ denote the resulting address space. Let π_\star denote the permutation on $(\mathfrak{S}_\star, \varpi_\star)$ that will achieve the same effect as a given permutation π on (\mathfrak{S}, ϖ) . Finally, let $\text{splits}_\star(i)$ denote the set of (one or two) digit positions in (\mathfrak{S}, ϖ) that digit position i becomes in $(\mathfrak{S}_\star, \varpi_\star)$.

⁵Note that there is no need to explicitly recombine decomposed digits, because we always store records compactly, and their order in a decomposed address space is consistent with the un-decomposed one.

Algorithm 6: Enumerate $(\mathfrak{S}, \varpi, \theta_Q, P)$.

Input: (\mathfrak{S}, ϖ) : specification of an n -digit permuted decomposed input address space; θ_Q : a setting for a component-wise prefix Q in this space; $P \supseteq Q$: a “bigger” component-wise prefix in the same space.
Output: the next valid setting for P compatible with θ_Q in ascending order, or signal the end of the enumeration if no such setting exists.

// initialize:
1 **foreach** $i \in Q$ **do** $\mathbf{v}[i] \leftarrow \theta_Q(i)$;
2 **foreach** $i \in P \setminus Q$ **do** $\mathbf{v}[i] \leftarrow 0$;
3 **foreach** $T \in \mathfrak{S}$ **do**
4 $\mathbf{p}[T] \leftarrow \min\{\varpi^{\text{pos}}(i) \mid i \in [p, n] \wedge \varpi^{\text{tree}}(i) = T\}$;
5 $\eta \leftarrow |\text{Leaves}(T)|$;
6 **foreach** $k \in [\mathbf{p}[T], \eta]$ **do** $\mathbf{z}[T, k] \leftarrow \kappa^T(\langle 0^{\eta-k}, \perp^k \rangle)$; // Lemma 4
7 **foreach** $k \in [\mathbf{p}[T], \eta]$ **do** $\mathbf{s}[T, k] \leftarrow \kappa^T(\Omega_{[k, \eta]}(\Omega_T(\mathbf{v})))$; // Lemma 5
// iterate:
8 **while true do**
9 $i \leftarrow \min(P)$;
10 **while** $i < \min(Q)$ **and** $\mathbf{v}[i] + 1 \geq \lceil \frac{\mathbf{s}[\varpi^{\text{tree}}(i), \varpi^{\text{pos}}(i+1)]}{\mathbf{z}[\varpi^{\text{tree}}(i), \varpi^{\text{pos}}(i)]} \rceil$ **do**
11 $i \leftarrow \min(P \cap [i+1, n])$;
12 **if** $i = \min(Q)$ **then return**; // signal the end of the enumeration
13 $\mathbf{v}[i] \leftarrow \mathbf{v}[i] + 1$;
14 $T \leftarrow \varpi^{\text{tree}}(i)$; $k \leftarrow \varpi^{\text{pos}}(i)$;
15 **if** $k > \mathbf{p}[T]$, $\mathbf{v}[i] = \lceil \frac{\mathbf{s}[T, k+1]}{\mathbf{z}[T, k]} \rceil - 1$, and $\mathbf{s}[T, k+1] \not\equiv 0 \pmod{\mathbf{z}[T, k]}$ **do**
16 **then** $\mathbf{s}[T, k] \leftarrow \mathbf{s}[T, k+1] \pmod{\mathbf{z}[T, k]}$;
17 **foreach** $j \in P \setminus [i, n]$ **do**
18 $\lfloor \mathbf{v}[j] \leftarrow 0$; $\mathbf{s}[\varpi^{\text{tree}}(j), \varpi^{\text{pos}}(j)] \leftarrow \mathbf{z}[\varpi^{\text{tree}}(j), \varpi^{\text{pos}}(j)]$;
19 **yield** \mathbf{v} ;

Algorithm 7: Optimize $_{\text{pass}}(\mathfrak{S}, \varpi, \pi)$: find the best one-pass plan, which may involve filtered reads.

Input: (\mathfrak{S}, ϖ) : an input address space; π : an ADP.
Output: a plan $(\mathfrak{S}_\star, \varpi_\star, \pi_\star, A, \beta)$ with the lowest cost.

1 $G_{\text{opt}} \leftarrow \perp$; $C_{\text{opt}} \leftarrow \infty$;
2 $\beta^{\min} \leftarrow \max\{i \mid \text{maxSize}^{\mathfrak{S}, \varpi}([0, i-1]) < B\}$;
3 $\beta^{\max} \leftarrow \max\{i \mid \text{maxSize}^{\mathfrak{S}, \varpi}([0, i-1]) \leq M/2\}$;
4 **for** $(\beta \leftarrow \beta^{\min}; \beta \leq \beta^{\max}; \beta \leftarrow \beta + 1)$ **do**
5 **foreach** $A_{\text{all-out}} \subseteq [0, \min\{\beta^{\min}, \beta - 1\}]$ **do**
6 **foreach** $\star \in \text{Decomp}(\mathfrak{S}, \varpi, \{\beta - 1\} \cup [0, \beta^{\min}] \setminus A_{\text{all-out}})$ **do**
7 $\beta' \leftarrow 1 + \min(\text{splits}_\star(\beta - 1))$;
8 $A \leftarrow \{\min(\text{splits}_\star(i)) \mid i \in [0, \beta - 1] \setminus A_{\text{all-out}}\}$;
9 **if not** $B \leq \text{maxSize}^{\mathfrak{S}_\star, \varpi_\star}(A) \leq M$ **then continue**;
10 $G \leftarrow \perp$;
11 **while true do**
12 $\hat{A} \leftarrow \pi_\star^{-1}([\phi(\pi_\star, A), \rho(\pi_\star, A)])$;
13 **if** $\hat{A} \cap [0, \beta'] \neq \emptyset$ or $\text{maxSize}^{\mathfrak{S}_\star, \varpi_\star}(A) > M$ **then break**;
14 $A \leftarrow A \cup \hat{A}$;
15 **if** $\text{mem}(\mathfrak{S}_\star, \varpi_\star, \pi_\star, A, \beta') \leq M$ **then**
16 $\lfloor G \leftarrow (\mathfrak{S}_\star, \varpi_\star, \pi_\star, A, \beta')$;
17 **if** $G \neq \perp$ and $\text{cost}(G) < C_{\text{opt}}$ **then**
18 $\lfloor G_{\text{opt}} \leftarrow G$; $C_{\text{opt}} \leftarrow \text{cost}(G_{\text{opt}})$;
19 **return** G_{opt} ;

Note that a digit with radix r has $\Theta(\sqrt{r})$ possible decompositions, namely: $(2, \lceil \frac{r}{2} \rceil)$, $(3, \lceil \frac{r}{3} \rceil)$, \dots , $(\lfloor \sqrt{r} \rfloor, \lceil \sqrt{r} \rceil)$, as well as their swapped versions $(\lceil \sqrt{r} \rceil, \lfloor \sqrt{r} \rfloor)$, \dots , $(\lceil \frac{r}{2} \rceil, 2)$.

Finding an Optimal Single-Pass Plan We begin with the task of finding an optimal single-pass plan for ADP_{pass} . Given (\mathfrak{S}, ϖ) and π , algorithm $\text{Optimize}_{\text{pass}}$ (shown as Algorithm 7) finds a plan $(\mathfrak{S}_\star, \varpi_\star, \pi_\star, A, \beta)$, where $(\mathfrak{S}_\star, \varpi_\star)$ represents an address space that is possibly further decomposed from (\mathfrak{S}, ϖ) , π_\star encodes the permutation in $(\mathfrak{S}_\star, \varpi_\star)$ equivalent to π in (\mathfrak{S}, ϖ) .

On a high level, $\text{Optimize}_{\text{pass}}$ searches through all viable input segment sizes between B and M . We use β to drive this search: $\text{maxSize}^{\mathfrak{S}, \varpi}([0, \beta])$ will be the segment size, and we will consider

decomposing the digit at position $\beta - 1$ in (\mathcal{X}, ϖ) to enable additional choices of segment sizes.

It can be shown (Lemma 11 in appendix) that an optimal plan includes in its action digits A all “out-digits”⁶ in $[0, \beta)$ fully (i.e., with no decomposition). Thus, we consider, as a choice for A^+ , every subset of in-digits or higher-order components⁷ from their decompositions. More specifically, in Algorithm 7, $A_{\text{all-out}}$ (Line 5) denotes the in-digits that fully belong to A^+ . The Decomp call (Line 6) considers all potential decompositions of $\beta - 1$ (to enable additional choices of segment sizes) as well as in-digits not already chosen by $A_{\text{all-out}}$ (to enable inclusion of decomposed digits in A^+). Then, we obtain a baseline choice for A (Line 8) consisting of the lower-order component of the decomposition of the digits $[0, \beta]$ not in $A_{\text{all-out}}$ (or the digit itself if it is not decomposed). This baseline choice includes all digits that must be in A given β' (which specifies the segment size in the decomposed address space) and A^+ .

Finally, we enumerate all choices for A given the baseline choice. This task turns out to be easy because the best choices can be enumerated by filling the gaps between action digits in the output address space, a complete gap at a time, from lower-order to higher-order (Lemma 12).

To sum up, $\text{Optimize}_{\text{pass}}$ always finds an optimal one-pass plan.

Theorem 3. *Optimize_{pass} finds the optimal one-pass plan for ADP_{pass} under the cost model of Eq. (6).*

Finding a Two-Pass Strategy To find a two-pass strategy, algorithm $\text{Optimize}_{2\text{pass}}$ (Algorithm 8) searches the space of possible intermediate address spaces between the two passes. Intuitively, performing a given ADP π entails moving those digits that are either “outgoing” (from digit positions below the input segment boundary to above the boundary) or “incoming” (from above the flush boundary to below the flush boundary). Let us denote these digits P and Q , respectively. $\text{Optimize}_{2\text{pass}}$ considers the possibility of decomposing digits in $P \cup Q$, so that the first pass can “partially” accomplish π by moving a subset of digits in $P \cup Q$ or their decomposed components to desired positions.

Since P by definition depends on the choice of segment size, $\text{Optimize}_{2\text{pass}}$ first considers each possible segment size in its outermost loop (as in $\text{Optimize}_{\text{pass}}$). We define P conservatively to be all “in-segment” digits: even if a digit stays “in-segment” after π , we do not rule out the possibility of moving it (or a component thereof) out temporarily in the first pass. Helper function Divide, defined using Decomp, enumerates all ways of decomposing a given set of digits and separating their higher-order and lower-order children. Line 6 calls Divide to consider ways to decompose the digits in P and move a subset P^\downarrow of the resulting digits “out” in the first pass.

The next step is to choose the set Q of candidate incoming digits. Lines 7–13 consider digits in the final address space that do not come from P ; they are added to Q one at a time in the order that they appear in the final address space (from low to high), until there is not enough memory to “hold” $P^\downarrow \cup Q$. The intuition behind the processing order and memory check is that 1) the actions digits for the second pass must include all of P^\downarrow (as they either stay in-segment or must move out-segment); 2) to get a big output flush in the second pass, we need a long consecutive sequence

⁶Here, a position i is *out-digit* if $\text{maxSize}^{\mathcal{X}, \varpi}([0, i]) \geq B$, or *in-digit* otherwise.

⁷For plans where A and A^+ contain components from the same original digit, we consider only those where A contains the lower-order component and A^+ contains the higher-order one. We do not consider any plan without this property, because it would be effectively the same cost-wise as a plan with this property, where the decomposition order is reversed for the digit.

of lowest-order action digits; and 3) the memory must be able to “hold” the action digits (i.e., all records compatible with a setting of non-action digits must fit in memory). Line 15 calls Divide to consider ways to further decompose the digits in Q so that a subset Q^\downarrow of the resulting digits can be moved “in” in the first pass.

Given the division of P into P^\downarrow (in-segment digits that stay in the first pass) and P^\uparrow (in-segment digits that move out in the first pass), as well as Q into Q^\downarrow (in-flush digits that enter in the first pass) and Q^\uparrow (in-flush digits that enter in the second pass), Line 20 considers two arrangements of the digits in the intermediate address space. Both arrangements place P^\downarrow at the lowest-order positions, followed by Q^\downarrow . However, they differ in where they place P^\uparrow , exposing a trade-off between read and write costs. The first arrangement places P^\uparrow next, followed by the remaining digits at the highest-order positions; hence, the first pass has a single write partition and low write cost. On the other hand, the second arrangement places P^\uparrow at the highest-order positions, pushing Q^\uparrow lower-order; hence, the second pass potentially has lower read cost.

Given an intermediate address space, Lines 22–25 determine the ADPs to be carried out in two steps, and optimize each by invoking $\text{Optimize}_{\text{pass}}$. The two-pass plan with the lowest total cost so far is remembered.⁸ Our implementation of $\text{Optimize}_{2\text{pass}}$ actually uses a simplified version of $\text{Optimize}_{\text{pass}}$ that does not consider filtered reads. Proving the optimality of $\text{Optimize}_{2\text{pass}}$ (for plans involving no filtered reads) remains an open problem, though empirically we observe that $\text{Optimize}_{2\text{pass}}$ produces very good plans from our experiments in Section 5.

Finding a Multi-Pass Strategy The search space becomes unwieldy for multi-pass plans, so we further restrict the search space and take a greedy approach. On a high level, $\text{Optimize}_{\text{multi-pass}}$ (Algorithm 9) greedily chooses a plan (without filtered reads) for the current pass that leads to a remaining permutation with the lowest “difficulty.” Function Estimate measures the “difficulty” of a permutation as the estimated amount of memory needed to perform the permutation in one pass without filtered reads—intuitively, more difficult permutations need more memory.

In more detail, when deciding how to decompose the original address space for its first step, $\text{Optimize}_{\text{multi-pass}}$ considers decomposing in-segment and in-flush digits (Lines 2 and 3, respectively), but it assumes that both segment and flush sizes are no less than B and close to B (within one digit). When deciding how to arrange digits in the intermediate address space resulted from the first step (Line 9), $\text{Optimize}_{\text{multi-pass}}$ simply places P^\downarrow (in-segment digits that stay in the first step) at the lowest-order positions, followed by Q^\downarrow (in-flush digits that enter in the first step), then P^\uparrow (in-segment digits that move out in the first step), and finally the remaining digits at the highest-order positions. This arrangement is basically the first one of the two options considered by the two-pass optimization algorithm $\text{Optimize}_{2\text{pass}}$. As discussed earlier, this option aims at saving write cost of the first step. The loop starting on Line 5 finds the intermediate address that leads to a remaining permutation with lowest estimated difficulty.

Overall Algorithm Finally, we come to ADP (Algorithm 10) for finding and executing a plan for a general ADP. The algorithm calls $\text{Optimize}_{\text{pass}}$, $\text{Optimize}_{2\text{pass}}$, and $\text{Optimize}_{\text{multi-pass}}$ to find the best one-, two-, and multi-pass plans, and picks the one with the lowest cost. Each pass is then executed using ADP_{pass} .

⁸Because of optimization by $\text{Optimize}_{\text{pass}}$, the actual segment and flush sizes and action digits used in the two passes may differ from those anticipated by $\text{Optimize}_{2\text{pass}}$. There is no problem with this discrepancy, as $\text{Optimize}_{2\text{pass}}$ ’s choices are used only to guide exploration of possible intermediate address spaces.

Algorithm 8: Optimize_{2pass}($\mathfrak{X}, \varpi, \pi$): find a two-pass plan.

Input: (\mathfrak{X}, ϖ): an input address space; π : an ADP.
Output: a pair of one-pass plans that together perform π .

```

1  $G_{\text{opt}} \leftarrow \perp$ ;  $C_{\text{opt}} \leftarrow \infty$ ;
2  $\beta^{\min} \leftarrow \max\{i \mid \text{maxSize}^{\mathfrak{X}, \varpi}([0, i-1]) < B\}$ ;
3  $\beta^{\max} \leftarrow \max\{i \mid \text{maxSize}^{\mathfrak{X}, \varpi}([0, i-1]) \leq M/2\}$ ;
4 for ( $\beta \leftarrow \beta^{\min}$ ;  $\beta \leq \beta^{\max}$ ;  $\beta \leftarrow \beta + 1$ ) do
5    $P \leftarrow [0, \beta)$ ; // "outgoing" digits; positions in space ( $\mathfrak{X}, \varpi$ )
6   foreach ( $\Delta, P^\uparrow, P^\downarrow \in \text{Divide}(\mathfrak{X}, \varpi, P)$ ) do // split outgoing
7      $Q \leftarrow \emptyset$ ; // find "incoming" digits; positions in space ( $\mathfrak{X}_\Delta, \varpi_\Delta$ )
8     while ( $[0, \|\mathfrak{X}_\Delta\|] \neq P^\uparrow \cup P^\downarrow \cup Q$ ) do
9        $q \leftarrow \arg \min_{i \in [0, \|\mathfrak{X}_\Delta\|] \setminus P^\uparrow \setminus P^\downarrow} \pi_\Delta(i)$ ;
10      if  $q \in P^\uparrow$  then // skip; case covered by another possible division of  $P$ 
11         $Q \leftarrow \emptyset$ ; break;
12      if  $\text{maxSize}^{\mathfrak{X}_\Delta, \varpi_\Delta}(P^\downarrow \cup Q \cup \{q\}) > M$  then break;
13       $Q \leftarrow Q \cup \{q\}$ ;
14      if  $Q = \emptyset$  then continue;
15      foreach ( $\nabla, Q^\uparrow, Q^\downarrow \in \text{Divide}(\mathfrak{X}_\Delta, \varpi_\Delta, Q)$ ) do // split incoming
16         $\star \leftarrow \nabla \circ \Delta$ ; // net effect of decompositions:  $\Delta$  followed by  $\nabla$ 
17         $P_\star^\uparrow \leftarrow$  positions of  $P^\uparrow$  in ( $\mathfrak{X}_\star, \varpi_\star$ ) after  $\nabla$ ;
18         $P_\star^\downarrow \leftarrow$  positions of  $P^\downarrow$  in ( $\mathfrak{X}_\star, \varpi_\star$ ) after  $\nabla$ ;
19         $R \leftarrow [0, \|\mathfrak{X}_\star\|] \setminus P_\star^\downarrow \setminus P_\star^\uparrow \setminus Q^\downarrow$ ;
20        // try two arrangements for the intermediate address space:
21         $\Pi' \leftarrow \left\{ \text{Arrange}(\{P_\star^\downarrow, Q^\downarrow, P_\star^\uparrow, R\}, \pi_\star), \right.$ 
22         $\quad \left. \text{Arrange}(\{P_\star^\downarrow, Q^\downarrow, R, P_\star^\uparrow\}, \pi_\star) \right\}$ ;
23        foreach  $\pi' \in \Pi'$  do
24          // The Optimizepass calls below use a simplified version that
25          // considers no filtered reads:
26           $G_1 \leftarrow \text{Optimize}_{\text{pass}}(\mathfrak{X}_\star, \varpi_\star, \pi')$ ;
27           $G_2 \leftarrow \text{Optimize}_{\text{pass}}(\mathfrak{X}_\star, \varpi_\star \circ (\pi')^{-1}, \pi_\star \circ (\pi')^{-1})$ ;
28          if  $\text{cost}(G_1) + \text{cost}(G_2) < C_{\text{opt}}$  then
29             $G_{\text{opt}} \leftarrow (G_1, G_2)$ ;  $C_{\text{opt}} \leftarrow \text{cost}(G_1) + \text{cost}(G_2)$ ;
30
31 return  $G_{\text{opt}}$ ;
32 def  $\text{Divide}(\mathfrak{X}, \varpi, S)$  begin
33   // Given address space ( $\mathfrak{X}, \varpi$ ), enumerate all ways of decomposing the digits
34   // in  $S$  and separating higher-order and lower-order children into  $S^\uparrow$  and
35   //  $S^\downarrow$ , respectively (if a digit is not decomposed, it can go into either  $S^\uparrow$  or
36   //  $S^\downarrow$ ).
37   foreach  $S_{\text{all-out}} \subseteq S$  do
38     foreach  $\star \in \text{Decomp}(\mathfrak{X}, \varpi, S \setminus S_{\text{all-out}})$  do
39        $S^\uparrow \leftarrow \bigcup_{i \in S_{\text{all-out}}} \text{splits}_\star(i)$ ;  $S^\downarrow \leftarrow \emptyset$ ;
40       foreach  $i \in S \setminus S_{\text{all-out}}$  do
41         if  $|\text{splits}_\star(i)| = 2$  then  $S^\uparrow \leftarrow S^\uparrow \cup \{\max(\text{splits}_\star(i))\}$ ;
42          $S^\downarrow \leftarrow S^\downarrow \cup \{\min(\text{splits}_\star(i))\}$ ;
43       yield ( $\star, S^\uparrow, S^\downarrow$ );
44   return; // signal the end of the enumeration
45
46 def  $\text{Arrange}(\{I_k\}, \pi)$  begin
47   // Given a partitioning  $\{I_k\}$  of digits (where partitions are listed from lowest
48   // to highest order) and a target permutation  $\pi$ , return a permutation that
49   // rearranges digits as follows: 1) any digit in  $I_{k'}$  is higher-order than any
50   // digit in  $I_k$  where  $k' > k$ ; 2) two digits in the same  $I_k$  are ordered
51   // according to  $\pi$ .
52   return  $i \mapsto |\{j \mid j \in I_k \wedge \pi(j) < \pi(i)\}| + \sum_{j < k} |I_j|$ , where  $i \in I_k$ ;

```

To recap, our approach to optimization focuses on common cases. As discussed at the end of Sections 3.1 and 3.2, ADP^{sr} takes very few passes in practice. Address decomposition, discussed in Section 4.2, allows ADP to retain this property. Therefore, we have developed Optimize_{pass} to find provably optimal single-pass plans (possibly involving filtered reads). Optimize_{2pass} finds two-pass plans (without filtered reads) by extensive exploration of the possible intermediate address spaces. Finally, Optimize_{multi-pass} uses a greedy heuristic to find multi-pass plans. As a result of this “tiered” optimization approach, the optimization overhead of ADP is very low in practice. For all experiments of Section 5, optimization took under one second.

Algorithm 9: Optimize_{multi-pass}($\mathfrak{X}, \varpi, \pi$): find a multi-pass plan.

Input: (\mathfrak{X}, ϖ): an input address space; π : an ADP.
Output: a sequence of one-pass plans that together perform π .
// The Optimize_{pass} calls below use a simplified version that considers no filtered reads.

```

1  $G \leftarrow \text{Optimize}_{\text{pass}}(\mathfrak{X}, \varpi, \pi)$ ; if  $G \neq \perp$  then return ( $G$ );
2  $P \leftarrow [0, \text{FindB}(\mathfrak{X}, \varpi))$ ;
3  $Q \leftarrow \pi^{-1}([0, \text{FindB}(\mathfrak{X}, \varpi \circ \pi^{-1})) \setminus P$ ;
4  $G_1 \leftarrow \perp$ ;  $C_{\text{opt}} \leftarrow \infty$ ;  $\star_1 \leftarrow \perp$ ;  $\pi_1 \leftarrow \perp$ ;
5 foreach ( $\star, U^\uparrow, U^\downarrow \in \text{Divide}(\mathfrak{X}, \varpi, P \cup Q)$ ) do
6    $P^\uparrow \leftarrow U^\uparrow \cap \bigcup_{i \in P} \text{splits}_\star(i)$ ;  $P^\downarrow \leftarrow U^\downarrow \cap \bigcup_{i \in P} \text{splits}_\star(i)$ ;
7    $Q^\uparrow \leftarrow U^\uparrow \cap \bigcup_{i \in Q} \text{splits}_\star(i)$ ;  $Q^\downarrow \leftarrow U^\downarrow \cap \bigcup_{i \in Q} \text{splits}_\star(i)$ ;
8    $R \leftarrow [0, \|\mathfrak{X}_\star\|] \setminus P^\downarrow \setminus P^\uparrow \setminus Q^\downarrow$ ;
9    $\pi' \leftarrow \text{Arrange}(\{P^\downarrow, Q^\downarrow, P^\uparrow, R\}, \pi_\star)$ ;
10   $G \leftarrow \text{Optimize}_{\text{pass}}(\mathfrak{X}_\star, \varpi_\star, \pi')$ ;
11   $C \leftarrow \text{Estimate}(\mathfrak{X}_\star, \varpi_\star \circ (\pi')^{-1}, \pi_\star \circ (\pi')^{-1})$ ;
12  if  $G \neq \perp$  and  $C < C_{\text{opt}}$  then
13     $G_1 \leftarrow G$ ;  $C_{\text{opt}} \leftarrow C$ ;  $\star_1 \leftarrow \star$ ;  $\pi_1 \leftarrow \pi'$ ;
14
15 return  $\text{concat}((G_1), \text{Optimize}_{\text{multi-pass}}(\mathfrak{X}_{\star_1}, \varpi_{\star_1} \circ \pi_1^{-1}, \pi_{\star_1} \circ \pi_1^{-1}))$ ;
16 def  $\text{FindB}(\mathfrak{X}, \varpi)$  begin
17   return  $\max\{i \mid \text{maxSize}^{\mathfrak{X}, \varpi}([0, i-1]) < B\}$ ;
18 def  $\text{DiffB}(\mathfrak{X}, \varpi)$  begin
19   return  $\text{maxSize}^{\mathfrak{X}, \varpi}([0, \text{FindB}(\mathfrak{X}, \varpi)]) - B$ ;
20 def  $\text{Estimate}(\mathfrak{X}, \varpi, \pi)$  begin
21   // Given ( $\mathfrak{X}, \varpi$ ), estimate the amount of memory needed to perform  $\pi$  in one
22   // pass without filtered reads, as a measure of  $\pi$ 's “difficulty.”
23    $U \leftarrow \{\text{FindB}(\mathfrak{X}, \varpi) - 1, \pi^{-1}(\text{FindB}(\mathfrak{X}, \varpi \circ \pi^{-1})) - 1\}$ ;
24    $\star \leftarrow \arg \min_{\star \in \text{Decomp}(\mathfrak{X}, \varpi, U) \wedge \text{DiffB}(\mathfrak{X}_\star, \varpi_\star) \geq 0 \wedge \text{DiffB}(\mathfrak{X}_\star, \varpi_\star \circ \pi_\star^{-1}) \geq 0}$ 
25      $(\text{DiffB}(\mathfrak{X}_\star, \varpi_\star) + \text{DiffB}(\mathfrak{X}_\star, \varpi_\star \circ \pi_\star^{-1}))$ ;
26    $\beta \leftarrow \text{FindB}(\mathfrak{X}_\star, \varpi_\star)$ ;  $\beta' \leftarrow \text{FindB}(\mathfrak{X}_\star, \varpi_\star \circ \pi_\star^{-1})$ ;
27    $A \leftarrow [0, \beta) \cup \pi_\star^{-1}([0, \beta'))$ ;
28   return  $\text{mem}(\mathfrak{X}_\star, \varpi_\star, \pi_\star, A, \beta)$ ;

```

Algorithm 10: ADP(R, τ, π): multi-pass, mixed-radix ADP.

Input: R : a file of records; τ : radices of the input address space; π : an ADP.
Output: a file containing permuted records.

```

1 initialize ( $\mathfrak{X}, \varpi$ ) from  $\tau$ ;
2  $\mathfrak{G} = \{\text{Optimize}_{\text{pass}}(\mathfrak{X}, \varpi, \pi), \text{Optimize}_{\text{multi-pass}}(\mathfrak{X}, \varpi, \pi)\}$ ;
3 if  $\mathfrak{G}$  contains a two-pass plan then  $\mathfrak{G} \leftarrow \mathfrak{G} \cup \{\text{Optimize}_{2\text{pass}}(\mathfrak{X}, \varpi, \pi)\}$ ;
4  $G \leftarrow \arg \min_{G \in \mathfrak{G}} \text{cost}(G)$ ;
5 foreach ( $\mathfrak{X}', \varpi', \pi', A', \beta'$ )  $\in G$  do
6    $R \leftarrow \text{ADP}_{\text{pass}}(R, \mathfrak{X}', \varpi', \pi', A', \beta')$ ;
7 return  $R$ ;


```

5 Experiments

We now evaluate ADP experimentally. Its main competitor will be a sorting-based implementation of permutation, which we refer to as SORT.⁹ SORT uses the external merge sort algorithm, with the following modifications. When producing intermediate runs, SORT writes out the input address of each record as a part of this record (so that these addresses can be used for comparison in subsequent passes). When producing the final result, SORT does not write these addresses. During merge phases, SORT fully utilizes the available memory to read and write in units as large as possible. For in-memory sorting, SORT uses an optimized quick-sort library. We implemented both ADP and SORT in C++.


We conduct experiments in two hardware settings. The first one is a single workstation (Intel i7 8-core 2.8GHz CPU, 8GB RAM, Linux kernel version 3.2.0) with several storage options: an Intel X25-E 32GB SSD, an Intel 320-S 80GB SSD, and a Samsung

⁹There are two other candidate competitors: Cormen’s BPC algorithm [4], and the brute-force strategy mentioned in Section 1. As discussed in Section 6, ADP is consistently better than BPC by a factor of at least 2, so the comparison would not be interesting. The brute-force strategy incurs too many reads and is one to two orders of magnitude slower than ADP for typical problems; see appendix for detailed results.



(a) **Figure 5:** $\mathcal{D}_{\text{TPC-H}}$ on SSD.

(b)



(a) **Figure 6:** $\mathcal{D}_{\text{TPC-H}}$ on cloud.

(b)

840 Pro 256GB SSD. To access the SSDs, we use an `ext4` file system without journaling but with direct I/O using the flag `O_DIRECT`. The block size is 4KB. Since experiments on all SSDs show similar results, we will only report results for 320-S here (see appendix for other results). In the following, “SSD” refers to this setting.¹⁰

The second hardware setting we consider is the “cloud.” Specifically, we use an `m1.medium` machine from Amazon EC2 to permute data stored in S3. We access S3 data through the `s3ql` file system, which stores files in the unit of blocks on S3. We set the block size to 64KB.

We use primarily two datasets. The first one, $\mathcal{D}_{\text{TPC-H}}$, derived from TPC-H [11], has following base dimensions:

<i>order-week</i>	<i>part-key</i>	<i>supplier-state</i>	<i>customer-state</i>	<i>order-priority</i>
574	100	50	50	2

The dimension sizes are chosen to reflect reality (e.g., 574 is the number of weeks during years 1995–2005). They are not perfect powers, and they provide a good mix of small and large dimensions. When needed, we vary the dataset size by scaling *part-key* and *supplier-state*. The second dataset, $\mathcal{D}_{v\text{-dim}}$, is designed to study the effect of the number of dimensions: given desired input size N (in the number of records) and number of dimensions d , we select d radices (each between $\lfloor \sqrt[d]{N} \rfloor$ and $\lceil \sqrt[d]{N} \rceil$) whose product is closest to N . The record size for both $\mathcal{D}_{\text{TPC-H}}$ and $\mathcal{D}_{v\text{-dim}}$ is 8 bytes.

$\mathcal{D}_{\text{TPC-H}}$ on SSD We begin by comparing ADP and SORT with a set of experiments for $\mathcal{D}_{\text{TPC-H}}$ on SSD. Figure 5a shows the running times when we vary the input size (by changing the size of the *part-key* dimension from 50 to 500 in increments of 50) while keeping memory at 5% of the input size; Figure 5b varies the amount of memory while keeping the input size fixed at 4GB. Each data point for ADP shows the average running time across five random permutations (SORT is oblivious to the permutation being performed).

From Figure 5a, we see that ADP’s lead over SORT improves as input size increases: for the smallest input, SORT is around 2.5 times slower than ADP, and for the largest input, it is around 3.6 times slower. Although SORT takes two passes throughout this figure, the culprit is the CPU time spent on sorting, which is still superlinear. In contrast, not only does ADP require only a single pass, it also spends no CPU time on comparisons.

From Figure 5b, we see similar advantages of ADP. For the two smallest memory sizes (460KB and 918KB), ADP uses 2 passes for four of the five randomly chosen permutations, and 1 pass for the last one, while SORT requires 4 and 3 passes respectively; on average, ADP is 2.7 and 2 times faster for these two memory sizes.

¹⁰Besides SSDs, we have also experimented with a traditional hard drive—a Seagate Barracuda 7200-RPM 250GB HD. While this setting is not what ADP targets, it helps illustrate, through comparison, how ADP exploits the characteristics of SSDs for efficiency; see appendix for details.




Figure 7: Percentage of permutations requiring specific numbers of passes. Darker shades mean more passes. For each input size, the left bar shows the breakdown for ADP; the right bar shows SORT.

Figure 8: Best single-pass strategies requiring specific numbers of passes vs. best two-pass strategy (when memory size varies, for each input size, the left bar shows the breakdown for ADP; the right bar shows SORT).

For all other memory sizes tested, ADP takes only 1 pass, whereas SORT takes 2. Overall, SORT is 2 to 3.3 times slower than ADP.

$\mathcal{D}_{\text{TPC-H}}$ on Cloud We run a similar set of experiments for $\mathcal{D}_{\text{TPC-H}}$ on the cloud. Figure 6a shows the running times when we vary the input size while keeping memory at 5% of the input size. Here, the advantage of ADP is similar to what we saw for the SSD setting: for the smallest input size, SORT is around 3.8 times to 2 times slower than ADP. The running times for both algorithms is markedly slower than on the SSD setting because S3’s throughput is much lower even with a block size that is 16 times larger. However, like in the SSD setting, there is little difference between random and sequential I/O for S3, so ADP shows similar gains. SORT’s running time increases in an almost linear fashion as we increase the input size, whereas ADP has some zig-zags because it chose different number of passes for some permutations.

Finally, we note that running time alone may not be the correct cost metric for this setting, because Amazon S3 actually charges by the amount of data stored and requested. Therefore, to get a clearer picture, we also show, in Figure 6b, the numbers of block I/Os performed for the same set of experiments in Figure 6a. Here, we again see a clear advantage of ADP over SORT.

All Permutations on $\mathcal{D}_{\text{TPC-H}}$ As the performance of ADP depends on the given permutation, here we study the behavior of ADP for all 119 non-identity permutations for $\mathcal{D}_{\text{TPC-H}}$. Figure 7 shows the percentage of the permutations requiring a specific number of passes. We fix the memory size at 128KB and vary the input size. We see that the majority of the permutations are “easy” (requiring 1 or 2 passes) for ADP, and they do not get much harder with larger inputs. In contrast, unable to exploit easy permutations, SORT always requires a high number of passes (invariant across permutations) compared with ADP, and even more with larger inputs.

Exploiting the Read/Write Cost Asymmetry Figure 8 compares the costs of the best single-pass strategy (with filtered reads enabled) and the best two-pass strategy, as we vary available memory. Here, $\alpha = 1/16$. In this case, it turns out that the minimum memory needed for a single-pass strategy without filtered reads is around 1.5MB; thus, for all memory sizes plotted in Figure 8, single-pass strategies are only possible with filtered reads. As we decrease memory, more aggressive use of filtered reads (i.e., with bigger waste factors) are needed to stay within a single pass. Thus, we see in Figure 8 that the read portion of the cost of the single-pass strategies continues to rise as memory gets smaller. Eventually (at the smallest memory size plotted in Figure 8), two-pass strategies become favorable again. The optimization built into ADP automatically picks the best strategy in all cases.

$\mathcal{D}_{v\text{-dim}}$ on SSD In this set of experiments, we fix the input size at 512MB and memory at 5% of the input size, and vary the number of dimensions d in $\mathcal{D}_{v\text{-dim}}$. Figure 9a compares the running times of ADP and SORT on SSD. We see that the number of dimen-




Figure 9: Effect of varying the number of dimensions d in $\mathcal{D}_{V\text{-dim}}$ on performance on SSD.

sions affects the SORT much more than ADP. ADP is around 2.1 times faster than SORT when $d = 2$ (this setting is of particular importance as it corresponds to matrix transpose, which has many applications). As d increases, this gap widens—SORT becomes progressively slower, while ADP remains efficient. When $d = 11$, ADP is around 3.1 times faster. Overall, ADP is around 2.1 to 3.4 times faster than SORT.

To further understand where the time is spent, we show in Figure 9b the CPU portion of the running times. While ADP has more complicated address calculation, SORT spends extra time in serializing/deserializing addresses as keys and comparing them. Overall, we see SORT’s CPU time rising linearly with d . We believe we can further widen this gap, since our current implementation of address calculation ADP leaves more room for improvement than the already optimized libraries used by SORT.

Another advantage of ADP is that it writes smaller intermediate results, while SORT has to write keys. This advantage is clearly visible in Figure 9c, which shows the amounts of data written by ADP and SORT. As d increases, SORT writes more because keys are larger. When $d = 11$, for example, SORT needs to perform 30% more block writes than it does when $d = 2$. Larger intermediate results can also lead to more passes for SORT, although Figure 9c does not show this behavior. On the other hand, ADP’s write size remains flat because ADP stores no keys; it is much lower than SORT also because ADP always takes one pass in this experiment.

6 Discussion and Related Work

Variable-Size Records We have so far assumed fixed-size records, which simplify the translation of linearized record addresses into locations within files. There are two methods to deal with variable-size records. The first method is to pad each record to its maximum size, which is simple and efficient if most records are near the maximum size. The second method is to add a level of indexing, e.g., B-tree, to enable translation of logical record addresses to physical locations. Instead of indexing each record, however, we only need to index each segment of records. Moreover, because our writes are appends in the granularity of blocks, we can index an intermediate result file efficiently like a bulk-loading procedure, thereby avoiding inefficiencies caused by random inserts (such as logarithmic insertion cost and half-empty blocks). During optimization, if record sizes are very skewed, we may use the average record size instead of the maximum size (which is more conservative) when deciding whether a set of action records can fit in memory. During execution, our algorithms can be adapted to handle unlucky situations when a large number of especially big action records overflow the available memory, but efficiency will suffer. Therefore, sorting may be more appealing in this case.

Sparse Data Section 4.2 shows how to deal with originally dense data that have *become* sparse because of address space decompositions. But what if data are already sparse in the original input address space, e.g., sparse matrices and data cubes? Analogous to the methods for handling variable-length records, there are two methods for handling sparse data. After all, one can think of a non-existent record as a record of size 0. The first method is to pad the

input address space, i.e., storing sparse data in a dense format. It only works when data are not too sparse. The second method is indexing, which works well if sparsity is random across the address space. Otherwise, sorting may be more appealing.

Related Work ADP is inspired by a class of well-known permutations called *BPC* (*bit-permute complement*). BPC is essentially a special case of ADP¹¹ where all radices are 2. In his dissertation [4], Cormen showed an asymptotically optimal algorithm for BPC in the parallel disk model. Each step of Cormen’s algorithm makes two passes: one to permute blocks into an intermediate order, and one to permute data within each memory worth of consecutive blocks. In comparison, our work targets storage solutions with better support for random accesses. By exploiting random reads more aggressively, we are able to complete such a step in a single pass, effectively reducing the number of I/Os by a factor of 2. We also consider optimizations such as filtered reads to further improve cost. Although our algorithms do not change the asymptotic complexity, we consider our constant factor improvement of 2 to be a very significant one, especially when it allows more problems to be solved in a single pass instead of two. As a special case of ADP (and BPC, if dimensions are powers of 2), matrix transpose has been studied extensively in the out-of-core setting [5, 6, 4, 10, 7]. Again, the target has been traditional hard drives.

Another major difference between our work and all previous work on BPC and matrix transpose is our handling of arbitrary radices and block/memory sizes. Almost all previous work conveniently assumes they are perfect powers of 2. Handling the general case requires padding, which, as we have argued in Section 1, will incur significant overhead in high dimensions. To the best of our knowledge, the only discussion of how to handle arbitrary radices efficiently is by Cormen [4] in the context of matrix transpose. He showed how, by carefully partitioning a matrix into 4 padded submatrices whose dimensions are powers of 2, the overall overhead of padding can be reduced. It is not clear how to generalize the technique to higher dimensions, where a small overhead in any two dimensions can still be blown up exponentially. We propose and evaluate a more general and robust solution.

As discussed in Example 2, one motivation for “resorting” data is computing multi-dimensional aggregates (e.g., [1, 14, 8]). This line of work mostly focused on optimizing the sharing of computation of multiple aggregates with one reorganization of data (e.g., using sorting, hashing, or traversal of a data array [14]). In general, multiple data reorganizations are still needed to compute all aggregates for a data cube—which is where our work can help in reducing the cost of the reorganizations themselves. Interestingly, Ross and Srivastava briefly mentioned [8] an optimization where sorting by dimensions (C, A, D) can benefit if data is already sorted by (C, D) , because we only need to further sort (consecutive) records with the same C value. Our work can be seen as a more general framework that systematically exploits such opportunities and other less obvious ones—such as using data sorted by (C, D) to produce data sorted by (D, C) more efficiently. Recent work on cooperative sorting [3] shares a similar motivation. The main idea there is to perform multiple sorts simultaneously in order to share work; the approach is still sorting-based and targets hard drives.

Finally, interest in *radix sort* and *radix join* continues to grow, especially on modern architectures (e.g., [9, 2]). Like our work, these algorithms operate on keys in radix-based representations. However, a fundamental difference is that for sort and join, the

¹¹BPC also allows an address bit to be complemented. The analogy in general ADP would be to allow addresses to be sorted by particular digits in descending instead of ascending order. This extension is straightforward and we omit it for brevity.

address or presence of an output record cannot be inferred from its key alone. Therefore, unlike our algorithms for ADP, they cannot escape the additional cost factor logarithmic in the input size. Nonetheless, recent work on efficiently implementing these algorithms may give insight on how to improve the CPU cost of ADP.

7 Conclusion

In this paper, we have introduced the notion of address-digit permutation (ADP), which captures a useful class of data reorganizations. We have shown how to perform ADP efficiently on modern block-based storage solutions that support random accesses better than traditional hard drives. By exploiting the characteristics of such storage, and by removing assumptions of previous work involving permutation, we have made ADP a better alternative to sorting in many practical settings.

An obvious direction of future work is more efficient implementations of ADP, especially its in-memory permutations on modern multicore architectures. Intuitively, ADP is perfectly suited to parallelization, because every record can be processed independently of all others. However, the effect of the memory hierarchy must be carefully considered. Another interesting problem for future work is how to combine the advantages of ADP and sorting. ADP works best when data are dense or uniformly sparse, while sorting handles skewed sparsity with ease. There may be a way to get the best of both approaches by judiciously choosing one of them for each subset of data or each step in an overall algorithm.

References

- [1] Sameet Agarwal, Rakesh Agrawal, Prasad Deshpande, Ashish Gupta, Jeffrey F. Naughton, Raghu Ramakrishnan, and Sunita Sarawagi. On the computation of multidimensional aggregates. In *Proceedings of the 1996 International Conference on Very Large Data Bases*, pages 506–521, Mumbai (Bombay), India, September 1996.
- [2] Cagri Balkesen, Jens Teubner, Gustavo Alonso, and M. Tamer Özsu. Main-memory hash joins on multi-core CPUs: Tuning to the underlying hardware. In *Proceedings of the 2013 International Conference on Data Engineering*, Brisbane, Australia, April 2013.
- [3] Yu Cao, Ramadhana Bramandia, Chee-Yong Chan, and Kian-Lee Tan. Optimized query evaluation using cooperative sorts. In *Proceedings of the 2010 International Conference on Data Engineering*, pages 601–612, Long Beach, California, USA, March 2010.
- [4] Thomas H. Cormen. *Virtual Memory for Data-Parallel Computing*. PhD thesis, Massachusetts Institute of Technology, 1993.
- [5] J. O. Eklundh. A fast computer method for matrix transposing. *IEEE Transactions on Computers*, 21(7):801–803, July 1972.
- [6] S. D. Kaushik, Chua-Huang Huang, John R. Johnson, Rodney W. Johnson, and P. Sadayappan. Efficient transposition algorithms for large matrices. In *Proceedings of the 1993 ACM/IEEE Supercomputing Conference*, pages 656–665, Portland, Oregon, USA, November 1993.
- [7] Sriram Krishnamoorthy, Gerald Baumgartner, Daniel Cociorva, Chi-Chung Lam, and P. Sadayappan. On efficient out-of-core matrix transposition. Technical report, School of Computer and Information Science, The Ohio State University, 2003.
- [8] Kenneth A. Ross and Divesh Srivastava. Fast computation of sparse datacubes. In *Proceedings of the 1997 International Conference on Very Large Data Bases*, pages 116–125, Athens, Greece, August 1997.
- [9] Nadathur Satish, Changkyu Kim, Jatin Chhugani, Anthony D. Nguyen, Victor W. Lee, Daehyun Kim, and Pradeep Dubey. Fast sort on CPUs and GPUs: A case for bandwidth oblivious SIMD sort. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*, pages 351–362, Indianapolis, Indiana, USA, June 2010.
- [10] Jinwoo Suh and Victor K. Prasanna. An efficient algorithm for out-of-core matrix transposition. *IEEE Transactions on Computers*, 51(4):420–438, 2002.

- [11] Transaction Processing Performance Council. The TPC benchmark H, 1993. <http://www.tpc.org/tpch/>.
- [12] Jeffrey Scott Vitter. External memory algorithms and data structures. *ACM Computing Surveys*, 33(2):209–271, 2001.
- [13] Yihong Zhao. *Performance Issues of Multi-Dimensional Data Analysis*. PhD thesis, University of Wisconsin at Madison, 1998.
- [14] Yihong Zhao, Prasad Deshpande, and Jeffrey F. Naughton. An array-based algorithm for simultaneous multidimensional aggregates. In *Proceedings of the 1997 ACM SIGMOD International Conference on Management of Data*, pages 159–170, Tucson, Arizona, USA, May 1997.

APPENDIX

A Proofs for Section 3.2

Lemma 8. *If the cost function $\text{cost}(\delta)$ is flat for $\delta \in [1, m - b]$ and non-decreasing for $\delta \geq 1$, then the optimal number of passes is no more than $\lceil \frac{\delta(\pi)}{m-b} \rceil$.*

Proof. Let $c_1 = \text{cost}(1) = \dots = \text{cost}(m - b)$, which lower-bounds the cost of one pass. Consider a baseline strategy that uses $\lceil \frac{\delta(\pi)}{m-b} \rceil$ passes, each admitting at most $m - b$ digits. The cost of this baseline strategy is $c_1 \lceil \frac{\delta(\pi)}{m-b} \rceil$. Any strategy with q passes has cost of at least $c_1 q$. Therefore, the optimal number of passes cannot be more than $\lceil \frac{\delta(\pi)}{m-b} \rceil$. \square

Lemma 9. *If the cost function $\text{cost}(\cdot)$ is convex, then given the number of passes q , there exists an optimal assignment $\Delta = \{p_1, \dots, p_q\}$ of the number of digits to admit in each pass, such that $|p_i - p_j| \leq 1$ for all i, j .*

Proof. Suppose there exist $p_i, p_j \in \Delta$ such that $p_i - p_j \geq 2$. Let Δ' denote the assignment where the i -th pass admits $p_i - 1$ digits, the j -th pass admits $p_j + 1$ digits, while all other passes admit the same number of digits as Δ . The difference in the costs of Δ and Δ' is $\text{cost}(p_i) + \text{cost}(p_j) - \text{cost}(p_i - 1) - \text{cost}(p_j + 1) \geq 0$ because $\text{cost}(\cdot)$ is convex (by Karamata's Inequality). \square

Proof of Theorem 1. Consider all strategies with exactly q steps. By Lemma 9, the numbers of digits admitted across passes differ by at most one. Suppose q_1 of the passes admit $\delta + 1$ each, and the remaining $q - q_1$ passes admit δ each. We have $q_1(\delta + 1) + (q - q_1)\delta = \delta(\pi)$. Solving for δ , we have $\delta = \frac{\delta(\pi) - q_1}{q}$. Because δ is an integer, $\delta(\pi) \equiv q_1 \pmod{q}$. Recall that $q_1 \leq q$. If $q_1 < q$, then $q_1 = \delta(\pi) \pmod{q}$. If $q_1 = q$, which only happens when q divides $\delta(\pi)$, the case is subsumed by that of $q_1 = 0$. Thus, $q_1 = \delta(\pi) \pmod{q}$, proving the choice of q_1 on Line 3 of Algorithm 4.

By Lemma 8, the optimal number of passes is no more than $\lceil \frac{\delta(\pi)}{m-b} \rceil$. The loop starting on Line 2 of Algorithm 4 ensures that all potentially optimal numbers of passes are considered. \square

B Proofs for Section 4.2

Proof of Lemma 2. We prove by induction over the structure of T . If T has only one node (i.e., no decomposition), the lemma trivially holds. Assume that the lemma holds for T' (with $\eta - 1$ leaves); we will show that it also holds for T (with η leaves) obtained by decomposing digit at position q in T' (say with radix r) into two digits at positions $q + 1$ and q in T (say with radices r'' and r'). By definition of decomposition, an address in T' becomes an address in T by replacing the value x at position q with values $x'' = \lfloor x/r' \rfloor$ and $x' = x \pmod{r'}$ at positions $q + 1$ and q . Note that $x = 0$ is replaced with $x'' = 0$ and $x' = 0$.

We are interested how the set X of valid addresses in T compatible with a given prefix setting $\theta_{[p,\eta]}$ are partitioned according to their settings for digit position $p-1$. For $p \geq q+3$ or $p \leq q$, the partitioning of X is the same as in T' , and it is easy to verify all claims of the lemma using our inductive hypothesis. We now consider the cases of $p = q+2$ and $p = q+1$.

For $p = q+2$, given $\theta_{[q+2,\eta]}$, the partitioning of X into $\{X_{x''}\}$ by the setting x'' for position $q+1$ can be refined into $\{X_{(x'',x')}\}$ by further partitioning by the setting x' for position q , where $x = x''r' + x'$ would be the setting for position q in T' . Let Z'_0 denote the set of addresses in T' compatible with $\langle 0^{\eta-q}, \perp^{q-1} \rangle$. The lemma's three claims can be shown as follows:

1. By the inductive hypothesis, the range of x is $[0, t)$ for some $t < r$. Thus, the range of $x'' = \lfloor x/r' \rfloor$ is also $[0, t')$ for some t' , where $0 < t' \leq r''$ because $x'' \leq x/r' < r/r' \leq r''$.
2. Every $X_{x''}$ except the last one contains exactly r' $X_{(x'',x')}$'s (for $x' = 0, 1, \dots, r'-1$). By the inductive hypothesis, all $X_{(x'',x')}$'s but the last one are identical to Z'_0 (modulo settings for digit positions higher than $p-1$; we omit this qualifier below for brevity). The last $X_{(x'',x')}$ belongs to the last $X_{x''}$, so all $X_{x''}$'s but the last one are identical. It is easy to further verify that unless $t = 1$, X_0 is identical to Z_0 because all of its $r'-1$ sub-partitions are identical to Z'_0 .
3. Consider any address in $X_{t''-1}$ (i.e., $x'' = t'' - 1$). Setting its digits in positions $[p-1, \eta)$ to 0 amounts to setting x in T' to a smaller value, and higher-order digits to 0, which is a valid address by the inductive hypothesis. This address is obviously in Z_0 .

For $p = q+1$, given $\theta_{[q+1,\eta]}$, X is partitioned into $\{X_{x'}\}$. Each $X_{x'}$ contains all valid addresses that are compatible with $\theta_{[q+1,\eta]}$ and have x' for position q . Suppose $\theta_{[q+1,\eta]}(q+1) = x''$. Consider the setting $\vartheta_{[q+1,\eta-1]}$ in T' that sets its $\eta - q - 2$ highest-order digits as the $\eta - q - 2$ highest-order digits in $\theta_{[q+1,\eta]}$. Each $X_{x'}$ thus corresponds to some partition in $\{X_x\}$ obtained by partitioning the valid addresses in T' compatible with $\vartheta_{[q+1,\eta-1]}$ according to their value x at digit position q —specifically, the partition with $x = x''r' + x'$. Note that in this case, Z_0 is identical to Z'_0 , the set of addresses in T' compatible with $\langle 0^{\eta-q}, \perp^{q-1} \rangle$. The lemma's three claims can be shown as follows:

1. By the inductive hypothesis, the range of x is $[0, t)$ for some $t < r$. Thus, the range of $x' = x \bmod r'$ is $[0, t')$ for some t' , where obviously $0 < t' \leq r'$.
2. Every $X_{x'}$ is some X_x . By the inductive hypothesis, all X_x 's but the last one are identical to Z'_0 and therefore to Z_0 . Obviously, only the last $X_{x'}$ can possibly be the last X_x . Therefore, all $X_{x'}$'s but the last one are identical to Z_0 .
3. Consider any address in $X_{t'-1}$ (i.e., $x' = t' - 1$). Setting its digits in positions $[p-1, \eta)$ to 0 amounts to setting x in T' to a smaller value, and higher-order digits to 0, which is valid address by the inductive hypothesis. This address is obviously in Z_0 . \square

Lemma 10. *Suppose T is a recursive decomposition of a single digit into an η -digit address space. For each $u \in \text{Leaves}(T)$ corresponding to digit position i in T , there exists at least one valid address with prefix $\langle 0^{\eta-i-1}, \text{rad}(u) - 1, \perp^i \rangle$.*

Proof. Follows from a simple inductive argument over the structure of T and the definition of decomposition. \square

Proof of Lemma 3. Follows directly from Lemma 4 and 5. \square

Proof of Lemma 4. We prove by induction over the structure of T . Obviously, if T has only one node (i.e., no decomposition), $\kappa^T(\langle 0 \rangle) = 1$. Assume that the lemma holds for T' (with $\eta - 1$ leaves); we will show that it also holds for T (with η leaves) obtained by decomposing digit at position q in T' into two digits at positions $q+1$ and q in T .

Consider digit position i in T . If $i \leq q-1$, the corresponding node u refers to digit position i in T' . By the inductive hypothesis, $z(u) = \kappa^{T'}(\langle 0^{\eta-1-i}, \perp^i \rangle)$. It is easy to see that the prefix setting $\langle 0^{\eta-1-i}, \perp^i \rangle$ (for $i \leq q-1$) in T' becomes $\langle 0^{\eta-1}, \perp^i \rangle$ in T . Therefore $z(u) = \kappa^T(\langle 0^{\eta-i}, \perp^i \rangle)$.

If $i \geq q+2$, the corresponding node u refers to digit position $i-1$ in T' . By the inductive hypothesis, $z(u) = \kappa^{T'}(\langle 0^{\eta-i}, \perp^{i-1} \rangle)$. It is easy to see that the prefix setting $\langle 0^{\eta-i}, \perp^{i-1} \rangle$ (for $i \geq q+2$) in T' becomes $\langle 0^{\eta-i}, \perp^i \rangle$ in T . Therefore $z(u) = \kappa^T(\langle 0^{\eta-i}, \perp^i \rangle)$.

For digit positions $q+1$ and q in T , the corresponding nodes u'' and u' are higher-order and lower-order children of node u that corresponds to digit position q . By the inductive hypothesis, $z(u) = \kappa^{T'}(\langle 0^{\eta-1-p}, \perp^p \rangle)$.

- First, consider $z(u')$. It is easy to see that the prefix setting $\langle 0^{\eta-1-p}, \perp^p \rangle$ in T' becomes $\langle 0^{\eta-p}, \perp^p \rangle$ in T . Therefore, $z(u') = z(u) = \kappa^T(\langle 0^{\eta-p}, \perp^p \rangle)$.
- Next, consider $z(u'')$. By definition of decomposition, $\kappa^T(\langle 0^{\eta-p-1}, \perp^{p+1} \rangle) = \sum_{x'=0}^{r'-1} \kappa^{T'}(\langle 0^{\eta-p-1}, x', \perp^p \rangle)$. Note that $\kappa^{T'}(\langle 0^{\eta-p-1}, 0, \perp^p \rangle) = z(u)$ by the argument under the previous bullet. To show that $z(u'') = z(u)r'$, we only need to show that all summands equal $z(u)$. Using a similar argument as in the proof of Lemma 2, we note that the sets of valid addresses in T with prefixes $\langle 0^{\eta-p-1}, x', \perp^p \rangle$ for different values of x' correspond to partitions $\{X_x\}$ of addresses in T' with prefix $\langle 0^{\eta-p-1}, \perp^p \rangle$ by their settings for digit position p . By Lemma 10, the last of the X_x 's contains addresses compatible with $\langle 0^{\eta-p-1}, \text{rad}(u) - 1, \perp^{p-1} \rangle$. With decomposition into T , $\text{rad}(u) - 1$ is mapped to a value greater than 0 for digit position $p+1$ in T' . Therefore, the sets of valid addresses in T with prefixes $\langle 0^{\eta-p-1}, x', \perp^p \rangle$ for different values of x' cannot correspond to the last partition in $\{X_x\}$. By Lemma 2, they are equal in size. \square

Proof of Lemma 5. It is easy to see that $\kappa^T(\langle \perp^\eta \rangle) = r$. To compute $\kappa^T(\langle x_{\eta-1}, \dots, x_{i+1}, x_i, \perp^i \rangle)$, we apply Lemma 2. The set X of addresses compatible with $\langle x_{\eta-1}, \dots, x_{i+1}, \perp^{i+1} \rangle$ can be partitioned according to their settings for digit position i ; all partitions except the last one are identical to Z_0 , the set of addresses compatible with $\langle 0^{\eta-i}, \perp^i \rangle$ (modulo digit positions $[i, \eta)$); the last partition is a subset of Z_0 (again modulo digit positions $[i, \eta)$). Therefore, all partitions of X except the last one have size z_i , and the last one has size no greater than z_i . The total size of X is s_{i+1} . Clearly, the formula for $\kappa^T(\langle x_{\eta-1}, \dots, x_{i+1}, x_i, \perp^i \rangle)$ computes the size of the partition with x_i for digit position i . \square

Proof of Theorem 2. Since data are dense in the original address space, the settings of the original digits in a valid address are independent. The theorem thus follows naturally. \square

Proof of Lemma 6. Equations for $\text{maxSize}(\cdot)$ and $\text{outSize}(\cdot)$ follow naturally from their definitions.

To compute $\text{offset}_{[p,n]}^{\xi, \varpi'}(\langle x_{n-1}, \dots, x_0 \rangle)$, we essentially count, among those addresses compatible with $\langle x_{n-1}, \dots, x_{p+1}, x_p, \perp^p \rangle$, how many precede $\langle x_{n-1}, \dots, x_0 \rangle$. The total number can be computed by summing up, for each $i \in [0, p)$, the number of addresses

compatible with $\langle x_{n-1}, \dots, x_{i+1}, y, \perp^i \rangle$, where $y \in [0, x_i]$. Using Lemma 2 and Theorem 2, we can show that the number of addresses compatible with $\langle x_{n-1}, \dots, x_{i+1}, y, \perp^i \rangle$ for a given $y \in [0, x_i]$ is the same as s_i , the number of addresses compatible with $\langle x_{n-1}, \dots, x_{i+1}, 0, \perp^i \rangle$. Hence, the grand total is $\sum_{i=0}^{p-1} x_i s_i$. \square

Proof of Lemma 7. Follows directly from Lemma 2 and the observation that the settings of the original digits in a valid address are independent (assuming data are dense in the original address space). \square

C Proofs for Section 4.3

Lemma 11. *There exists an optimal one-pass plan that includes in its actions digits all digits in $[\beta_{\min}, \beta)$ fully (i.e., with no decomposition), where $\beta_{\min} = \max\{i \mid \max\text{Size}^{\mathbb{Z}, \infty}([0, i-1]) < B\}$.*

Proof. Suppose a plan \mathcal{P} whose action digits A does not include all digits in $[\beta_{\min}, \beta)$ fully. Consider an alternative plan \mathcal{P}' with the same set of action digits A , but simply sets its β instead to $\beta' = \beta_{\min}$. Recall the cost model in Eq. (6). Going from \mathcal{P} to \mathcal{P}' eliminates at least one non-action digit (or component thereof) and therefore reduce the waste factor term in C_r by at least by 1/2. On the other hand, since all plans have input segment sizes no less than B , $0 \leq U_r \leq \frac{N}{B}$, so going from \mathcal{P} to \mathcal{P}' cannot cause the term $(\frac{N}{B} + U_r)$ to increase by more than by factor of 2. Thus, C_r for \mathcal{P}' is at most that of \mathcal{P} . Furthermore, C_w remains unchanged because \mathcal{P}' does not change A . Overall, the cost of \mathcal{P}' is no greater than that of \mathcal{P} . Finally, the memory requirement stays the same, because both A and therefore the number of partitions per group are unchanged. \square

Lemma 12. *Consider all plans with given β and with $A \supseteq A_{\text{base}}$ for a given A_{base} for performing a given ADP π on an n -digit decomposed address space. There exists an optimal plan with the following property: for every digit $i \notin A$ and $\pi(i) > 0$, $[\pi(i) - 1, n) \cap \pi(A \setminus A_{\text{base}}) = \emptyset$. Informally, this plan completely fills zero or more “gaps” between A_{base} digits in the output address space, from lower-order to higher-order, by including all digits in these gaps in A .*

Proof. Let \mathcal{P} be a plan with $A \supseteq A_{\text{base}}$, where there exists $i \notin A$ and $\pi(i) > 0$ such that $[\pi(i) - 1, n) \cap \pi(A \setminus A_{\text{base}}) \neq \emptyset$. Let $j_1 = \min\{j \in [0, \pi(i)) \mid [j, \pi(i)) \cap \pi(A_{\text{base}}) = \emptyset\}$ and $j_2 = \max\{j \in (\pi(i), n] \mid [\pi(i), j) \cap \pi(A_{\text{base}}) = \emptyset\}$; i.e., $[j_1, j_2)$ is the gap containing $\pi(i)$ between A_{base} digits in the output address. We consider an alternative plan \mathcal{P}' with actions bits $A' = A \setminus \pi^{-1}([j_1, n)) \cup A_{\text{base}}$. We claim that $\rho(\pi, A) \leq \rho(\pi, A')$. To show this claim, there are two cases. First, if $\rho(\pi, A) < j_1$, then $\rho(\pi, A') = \rho(\pi, A)$ because $\pi(A)$ and $\pi(A')$ are identical over $[0, j_1)$. Second, if $\rho(\pi, A) \geq j_1$, it must be the case that $\rho(\pi, A) \leq j_2$, because at one digit in $[j_1, j_2)$ is not in $\pi(A)$ (namely, $\pi(i)$). On the other hand, $\rho(\pi, A') = j_2$. Therefore, $\rho(\pi, A) \leq \rho(\pi, A')$ for both cases, so from \mathcal{P} to \mathcal{P}' , the output partition size becomes bigger or remains the same.

Recall the cost model in Eq. (6). Going from \mathcal{P} to \mathcal{P}' does not change the input segment size or waste factor, so C_r stays the same; the output partition size is no less than before, so C_w is no more than before. Overall, the cost of \mathcal{P}' is no greater than that of \mathcal{P} . Finally, the memory requirement is lower, because A is smaller, and because the number of partitions per group is no more than before—as both the number of partitions overall and the number of action digits at or above the partition position are no more than before. \square




Figure 10: $\mathcal{D}_{\text{TPC-H}}$ on Samsung 840 Pro SSD.




Figure 11: $\mathcal{D}_{\text{TPC-H}}$ on Intel X25-E SSD.

Proof of Theorem 3. Follow directly from Lemmas 11 and 12, and the discussion of $\text{Optimize}_{\text{pass}}$ in Section 4.3. \square

D Additional Experimental Results

$\mathcal{D}_{\text{TPC-H}}$ on Samsung 840 Pro and Intel X25-E SSDs We compare ADP and SORT on Samsung 840 Pro and Intel X25-E SSDs, using the same experiments on $\mathcal{D}_{\text{TPC-H}}$ as in Figure 5 (which ran on Intel 320-S). Figure 10a and Figure 11a show the running times when we vary the input size while keeping memory at 5% of the input size; Figure 10b and Figure 11b vary the amount of memory while keeping the input size fixed at 4GB. Again, each data point shows the average running time across five random permutations. These figures show the same trends as those observed in Figure 5 for Intel 320-S, demonstrating that ADP offers consistent and substantial advantages over SORT across SSDs with different models and ages—before running the experiments, the amounts of data written were 3.7TB for Samsung 840 Pro, 9.0TB for Intel 320-S, and 14.7TB for Intel X25-E.

$\mathcal{D}_{\text{TPC-H}}$ on Seagate Barracuda 7200RPM HD Although ADP is not designed for hard disks, we have nonetheless experimented with ADP and SORT on a traditional hard disk (Seagate Barracuda 7200RPM), using the same experiments on $\mathcal{D}_{\text{TPC-H}}$ as in Figure 5. As mentioned earlier in Section 5, these experiments, through contrast with results on SSDs, help illustrate how ADP exploits the characteristics of SSDs for efficiency.

Because of ADP’s random accesses, it is natural to expect a pass in ADP to be much slower than a pass in SORT on a hard drive. However, for many permutations, this issue is ameliorated because ADP can read segments and write partitions that are much longer than a block, and it is designed to take advantage of this opportunity as much as possible. Thus, the comparison between ADP and SORT on hard drives is more interesting than one might expect.

Figure 12 shows the results. As in other experiments on $\mathcal{D}_{\text{TPC-H}}$, we use five random permutations. For SORT, performance is stable across permutations, so we only show its average running times. For ADP, we plot running times for all permutations in addition to the average trend. We see that performance of ADP is highly variable depending on the permutation—some permutations take longer than SORT because of random accesses, while others are faster thanks to fewer passes and bigger segments and partitions. In Figure 12a, we increase the size of *part-key* dimension from 10 to 100 in increments of 10, while keeping memory at 5% of the input size. As the figure shows, when we scale up both data and memory, ADP and SORT have comparable *average* performance, but ADP has higher variance. On the other hand, Figure 12b shows




Figure 12: $\mathcal{D}_{\text{TPC-H}}$ on Seagate Barracuda 7200-RPM hard disk.




Figure 13: Brute-force (naive) vs. ADP for $\mathcal{D}_{\text{TPC-H}}$ on Samsung 840 Pro.

that when memory is limited, ADP still beats SORT even on traditional hard drives, because it requires much fewer passes. Comparing with similar figures for SSDs, we see that traditional hard drives’ poor random access performance diminishes ADP’s benefit significantly, but ADP still outperforms SORT for many permutations or when memory is limited.

ADP vs. Brute-Force, $\mathcal{D}_{\text{TPC-H}}$ on Samsung 840 Pro We have also implemented the brute-force strategy described in Section 1 and compared it with ADP for $\mathcal{D}_{\text{TPC-H}}$ on Samsung 840 Pro. Figure 13 shows their running times as we vary the input size while keeping memory at 5% of the input size. Brute-force uses its available memory as an LRU block cache for reading input; random reads satisfied by the cache do not incur I/Os. Brute-force turned out very slow, so we tested smaller inputs: we set the size of the *supplier-state* dimension to 10 instead of 50, and we increased the size of the *part-key* dimension from 2 to 20 in increments of 2.

As Figure 13 shows, the brute-force strategy is up to two orders of magnitude worse than ADP, because of its large number of random record reads. For the largest input size tested here (88MB), brute-force on average takes 546 seconds, while ADP takes only 5.8 seconds, a 94-fold improvement. Overall, ADP is around 35 to 117 times faster than brute-force.

Figure 13 also shows that the performance of brute-force heavily depends on the permutation being performed. On rare occasions, brute-force can be as fast as ADP, when the output order happens to make the LRU cache extremely effective. For example, for the input size of 8MB, one permutation, (0, 3, 2, 1, 4), has high “locality”—it mostly involves permuting records within a small number of blocks. Thus, brute-force works well in this case (and so is ADP). The majority of the permutations do not have such locality, however, so brute-force is generally much slower than ADP even though SSDs support fast random reads.