

On the Value of Multiple Read/Write Streams for Approximating Frequency Moments

Paul Beame*

Computer Science and Engineering
University of Washington
Seattle, WA 98195-2350
beame@cs.washington.edu

Dang-Trinh Huynh-Ngoc†

Computer Science and Engineering
University of Washington
Seattle, WA 98195-2350
trinh@cs.washington.edu

Abstract

We consider the read/write streams model, an extension of the standard data stream model in which an algorithm can create and manipulate multiple read/write streams in addition to its input data stream. We show that any randomized read/write stream algorithm with a fixed number of streams and a sublogarithmic number of passes that produces a constant factor approximation of the k -th frequency moment F_k of an input sequence of length of at most N from $\{1, \dots, N\}$ requires space $\Omega(N^{1-4/k-\delta})$ for any $\delta > 0$. For comparison, it is known that with a single read-only data stream there is a randomized constant-factor approximation for F_k using $\tilde{O}(N^{1-2/k})$ space and that there is a deterministic algorithm computing F_k exactly using 3 read/write streams, $O(\log N)$ passes, and $O(\log N)$ space. Therefore, although the ability to manipulate multiple read/write streams can add substantial power to the data stream model, with a sub-logarithmic number of passes this does not significantly improve the ability to approximate higher frequency moments efficiently. Our lower bounds also apply to $(1 + \epsilon)$ -approximations of F_k for $\epsilon \geq 1/N$.

1. Introduction

The development of efficient on-line algorithms for computing various statistics on streams of data has been a remarkable success for both theory and practice. The main model has been the data stream model in which algorithms with limited storage access the input data in one pass as it streams by. This model is natural for representing many problems in monitoring web and transactional traffic.

However, as Grohe and Schweikardt [9] observed, in many natural situations for which the data stream model has been studied, the computation also has access to aux-

iliary external memory for storing intermediate results. In this situation, the lower bounds for the data stream model no longer apply. This motivated Grohe and Schweikardt to introduce a model, termed the *read/write streams model* in [5], to capture this additional capability. In the read/write streams model, in addition to the input data stream, the computation can manipulate multiple sequentially-accessed read/write streams.

As noted in [9], the read/write streams model is substantially more powerful than the ordinary data stream model since read/write stream algorithms can sort lists of size N with $O(\log N)$ passes and space using 3 streams. Unfortunately, given the large values of N involved, $\Theta(\log N)$ passes through the data is a very large cost. For sorting, lower bounds given in [9, 8] show that such small space read/write stream algorithms are not possible using fewer passes; moreover, [8, 5] show lower bounds for the related problems of determining whether two sets are equal and of determining whether or not the input stream consists of distinct elements.

While these lower bounds give us significant understanding of the read/write streams model they apply to exact computation and do not say much about the potential of the additional read/write streams model for more efficient solutions of approximation problems, which are the major successes of the standard data stream model (see surveys [12, 2]). Among the most notable successes are the surprising one-pass small space randomized algorithms for approximating the *frequency moments* of data streams first shown by Alon, Matias, and Szegedy [1]. The k -th frequency moment, F_k , is the sum of the k -th powers of the frequencies with which elements occur in a data stream. F_1 is simply the length of the data stream; F_0 is the number of distinct elements in the stream; if the stream represents keys of a database relation then F_2 represents the size of the self-join on that key. The methods in [1] also yield efficient randomized algorithms for approximating F_∞^* , the largest frequency of any element

* Research supported by NSF grants CCF-0514870 and CCF-0830626

† Research supported by a Vietnam Education Foundation Fellowship

in the stream. The current best one-pass algorithms for frequency moments approximate F_k within a $(1 + \epsilon)$ factor on streams of length N using $\tilde{O}(N^{1-2/k})$ space [11, 6].

Alon, Matias, and Szegedy also showed that their algorithms were not far from optimal in the data stream model; in particular, by extending bounds [13] for the randomized 2-party communication complexity for a promise version of the set disjointness problem from 2 to p players, where each player has access to its own private portion of the input (a model known as the *p-party number-in-hand communication game*) they showed that F_k requires $\Omega(N^{1-5/k})$ space to approximate by randomized algorithms. A series of papers [14, 3, 7] has improved the space lower bound to an essentially optimal $\tilde{\Omega}(N^{1-2/k})$ by improving the lower bound for the communication game; thus F_k for $k > 2$ requires polynomial space in the data stream model¹.

This leads to the natural questions: Can one prove good lower bounds for approximation problems in the read/write streams model? Can read/write stream algorithms approximate larger frequency moments more efficiently than one-pass algorithms can?

We show that the ability to augment the data stream model with multiple read/write streams does not produce significant additional efficiency in approximating frequency moments. In particular, any randomized read/write stream algorithm with a fixed number of streams and $o(\log N)$ passes that approximates the k -th frequency moment F_k of an input sequence of length of at most N from $\{1, \dots, N\}$ within a constant factor requires space $\Omega(N^{1-4/k-\delta})$ for any constant $\delta > 0$. This lower bound is very similar to the upper bound even for standard single pass read-only data streams (and is larger than the original lower bound in [1] for ordinary data streams).

The major difficulty in developing lower bounds for the read/write streams model, in contrast to the data streams model, is that an easy reduction from number-in-hand multiparty communication complexity breaks down. This fails for read/write stream algorithms because different parts of the computations can communicate with each other by writing to the streams. In fact, the p -party disjointness problem, which is the basis for the lower bounds for approximating frequency moments in the data streams model, can be easily solved by read/write stream algorithms using only 3 passes, 2 streams and $O(\log N)$ space.

The amount of data written on the streams also prevents the use of traditional time-space tradeoff lower bound methods, which are the other obvious tools to consider. As a result, previous work on lower bounds in the read/write streams model has been based on special-purpose combi-

natorial methods developed especially for the model.

Grohe, Hernich, and Schweikardt [9, 8] identified certain structural properties of the executions of read/write stream algorithms, their *skeletons*, and applied cut-and-paste arguments along with these skeletons to show the existence of certain combinatorial rectangles on which the algorithms' answers must be constant. They showed that the existence of these rectangles implies that no space-efficient read/write stream algorithm can sort in $o(\log N)$ passes or determine, with one-sided error bounded below $1/2$, whether or not two input sets are equal. Then, by reduction, they derived lower bounds for one-sided error randomized algorithms for several other problems.

Beame, Jayram, and Rudra [5] used the structure of the rectangles produced in [9, 8] together with additional combinatorial reasoning to show how standard properties that lower bound randomized two-party communication complexity – discrepancy and corruption over rectangles – can be used to derive lower bounds for randomized read/write streams with two-sided error. Using this approach they gave general methods for obtaining lower bounds for two-sided error randomized read/write stream algorithms. In particular they showed that with $o(\log N / \log \log N)$ passes and $O(N^{1-\delta})$ space, randomized read/write stream algorithms with two-sided error cannot determine whether or not two input sets are disjoint. This yielded several other lower bounds, including an $\Omega(N^{1-\delta})$ lower bound on the space for computing a 2-approximation of F_∞^* with $o(\log N / \log \log N)$ passes and a similar lower bound for exact computation of F_0 .

However, the methods of [9, 8, 5] do not yield lower bounds on the approximate computation of frequency moments F_k for any $k < \infty$. In particular it is consistent with all previous work that read/write stream algorithms can compute constant factor approximations to any such F_k using $o(\log N)$ passes, $O(\log N)$ space, and only 2 streams. We show that this is not possible.

We take a different approach to lower bounds in the read/write streams model from those in [9, 8, 5]. Despite the failure of the standard reduction, we are able to characterize read/write stream algorithms via a direct simulation of read/write stream algorithms by p -party communication protocols. Though quite different in the overall structure of the argument, this reduction does make use of a simplified variant of the skeletons defined in [9, 8]. Our method may have many other applications.

For the specific case of approximating frequency moments we derive our lower bounds by applying our simulation to a blocked and permuted version of the promise p -party disjointness problem (with p depending on N and k). The problem is a generalization of one considered in [5] extended to the case of p sets. This allows us to obtain the $\Omega(N^{1-4/k-\delta})$ space lower bounds for computing F_k using

¹The $\Omega(1/\epsilon^2)$ dependence of the space on the error in the approximation has also been shown to be optimal for one-pass algorithms [10, 16], though by considering a different two-party communication problem – approximate Hamming distance.

a sublogarithmic number of passes and a constant number of streams.

Although this nearly matches the best lower bounds for the data stream model, there is a gap between our read/write streams lower bounds and the data stream upper bounds; our lower bounds are limited by the relationship between the number of blocks and the number of sets in the permuted disjointness problem that we consider. We also show that modifying this relationship cannot improve the lower bound for constant factor approximations for $k < 3.5$. In particular, there is a deterministic read/write stream algorithm with three passes, two streams and $O(\log N)$ space that can compute the value of the blocked and permuted p -party disjointness problem for any numbers of blocks and sets that would have produced such lower bounds. To derive this algorithm we show a novel property on the lengths of common subsequences in sets of permutations.

2. Preliminaries

In the read/write streams model, the streams are represented as t read/write Turing machine tapes. The input stream is given as the contents of the first such tape; the other streams/tapes are used for working storage. Passes through the data in a stream correspond to reversals on the corresponding Turing machine tape; the number of passes is one more than the number of reversals. The internal memory of read/write streams allows random access.

The three resource parameters that are important to a read/write stream algorithm A are (1) the number of external read/write tapes t that A uses, (2) the maximum space s that A uses, and (3) the maximum number of reversals r made by A on all the external tapes.

Since we will primarily focus on lower bounds, we define a nonuniform version of the read/write stream model since lower bounds for this model are more general than those that only apply to the uniform case. Fix an input alphabet Σ and tape alphabet Γ . An (r, s, t) -read/write stream algorithm A on Σ^N is an automaton with 2^s states with one read/write head on each of t tapes. It begins with its input $\mathbf{v} \in \Sigma^N$ on the first tape and the remaining tapes blank. In each step, based on the current state and currently scanned symbols, one of its heads writes a new symbol from Γ in its currently scanned tape cell and moves one cell left or right. On any input $\mathbf{v} \in \Sigma^N$ it reverses the direction of movement of its heads at most r times before it halts.

For functions $r, s : \mathbb{N} \rightarrow \mathbb{N}$ and $t \in \mathbb{N}$, a (nonuniform) $(r(\cdot), s(\cdot), t)$ -read/write stream algorithm on Σ^* is a family of algorithms $\{A_N\}_{N \in \mathbb{N}}$ where for each N , A_N is an $(r(N), s(N), t)$ -read/write stream algorithm and all A_N have the same input and tape alphabets. Randomized and nondeterministic algorithms are defined analogously.

For integer $m \geq 1$ denote $\{1, \dots, m\}$ by $[m]$ and for any permutation ϕ of $[m]$, define the *sortedness* of ϕ , denoted

by $\text{sortedness}(\phi)$, to be the length of the longest monotone subsequence of $(\phi(1), \dots, \phi(m))$. For any $m, p \geq 1$ and a sequence $\Phi = (\phi_1, \phi_2, \dots, \phi_p)$ of permutations of $[m]$, define the *relative sortedness* of Φ , denoted by $\text{relsorted}(\Phi)$, to be $\max_{i \neq j \in [p]} (\text{sortedness}(\phi_i \phi_j^{-1}))$.

Lemma 2.1. *If ϕ is a random permutation of $[m]$ then $\Pr[\text{sortedness}(\phi) \geq 2e\sqrt{m}] < 2^{-4e\sqrt{m}+1}$.*

Proof. $\Pr[\text{sortedness}(\phi) \geq t] \leq 2\binom{m}{t}/t! < 2(me^2/t^2)^t. \square$

Corollary 2.2. *If $p = m^c$ for some constant $c > 0$ and m is sufficiently large there exists a sequence $\Phi = (\phi_1, \phi_2, \dots, \phi_p)$ of permutations of $[m]$ such that $\text{relsorted}(\Phi) < 2e\sqrt{m}$.*

For a given p polynomial in m , we let Φ^* be a sequence that is guaranteed to exist by Corollary 2.2.

The k -th frequency moment F_k of a sequence $a_1, \dots, a_n \in [m]$ is $\sum_{j \in [m]} f_j^k$ where $f_j = \#\{i \mid a_i = j\}$. We will typically consider the problem when $m = n$. Also write F_∞^* for $\max_{j \in [m]} f_j$.

For $2 \leq p < n$, define the promise problem $\text{PDISJ}_{n,p} : \{0, 1\}^{np} \rightarrow \{0, 1\}$ as follows: For $x_1, \dots, x_p \in \{0, 1\}^n$, interpret each x_i as the characteristic function of a subset of $[n]$. If these subsets are pair-wise disjoint then $\text{PDISJ}(x_1, \dots, x_p) = 0$; if there is a unique element $a \in [n]$ such that $x_i \cap x_j = a$ for all $i, j \in [p]$ then $\text{PDISJ}(x_1, \dots, x_p) = 1$; otherwise, PDISJ is undefined.

We use the usual definition of p -party number-in-hand communication complexity. A series of communication complexity lower bounds [1, 14, 3, 7] for $\text{PDISJ}_{n,p}$ in this number-in-hand model has resulted in essentially optimal lower bounds for computing frequency moments in the data stream model, even allowing multiple passes on the input stream. The strongest of these bounds [7] shows that any p -party public-coin randomized number-in-hand communication protocol for $\text{PDISJ}_{n,p}$ must have complexity at least $\Omega(\frac{n}{p \log p})$. (The bound is an even stronger $\Omega(n/p)$ for one-way communication.)

However, as noted in the introduction, for any n and p , there is a simple $(2, \log_2 n + O(1), 2)$ read/write stream algorithm for computing $\text{PDISJ}_{n,p}$: Copy x_1 to the second tape and compare the contents of x_1 and x_2 bit-by-bit using the two tape heads. We therefore will need a modified function in order to obtain our lower bounds for approximating frequency moments.

Let $N > p \geq 2$ and let $\Pi = (\pi_1, \dots, \pi_p)$ be a sequence of permutations on $[N]$. We define the promise problem $\text{PDISJ}_{N,p}^\Pi : \{0, 1\}^{Np} \rightarrow \{0, 1\}$ by $\text{PDISJ}_{N,p}^\Pi(y_1, \dots, y_p) = \text{PDISJ}_{N,p}(x_1, \dots, x_p)$ where the j -th bit of x_i is the $\pi_i^{-1}(j)$ -th bit of y_i . The relationship between x_i and y_i is equivalent to requiring that $y_{ij} = x_{i\pi_i(j)}$.

We first observe that the same reduction idea given by [1] yields lower bounds for F_k given lower bounds for $\text{PDISJ}_{N,p}^\Pi$ for suitable choices of p .

Lemma 2.3. Let $N > 1$ and $1/2 > \delta \geq 0$.

- (a) Let $k > 1$, $1 \geq \epsilon > 4^{1/(k-1)}/N$, and $p \geq (4\epsilon N)^{1/k}$. If there is a randomized (r, s, t) -read/write stream algorithm that outputs an approximation of F_k within a $(1 + \epsilon)$ -factor on a sequence of at most N elements in the range $[N]$ with probability at least $1 - \delta$ then there is a randomized $(r + 2, s + O(\log N), \max(t, 2))$ -read/write stream algorithm with error at most δ that solves $\text{PDISJ}_{N,p}^{\Pi}$ for any Π .
- (b) Let $k < 1$, $1/2 > \epsilon \geq 0$, and let $p - p^k \geq 2\epsilon N$. If there is a randomized (r, s, t) -read/write stream algorithm that outputs an approximation of F_k within a $(1 + \epsilon)$ -factor on a sequence of at most N elements in the range $[N]$ with probability at least $1 - \delta$ then there is a randomized $(r + 2, s + O(\log N), \max(t, 2))$ -read/write stream algorithm with error at most δ that solves $\text{PDISJ}_{N,p}^{\Pi}$ for any Π .

Proof. We only prove part (a); the proof for part (b) is similar. Given an input \mathbf{v} to $\text{PDISJ}_{N,p}^{\Pi}$, the read/write stream algorithm reads its bits from left to right to convert it on the second tape to the appropriate input \mathbf{v}' for the F_k problem in the obvious way: Read \mathbf{v} from left to right, for every bit that is as 1 in position $i \in [N]$ in the input, write i to \mathbf{v}' . While doing this also use the state to record the number $N' \leq N$ of 1's in the input. If $N' > N$ then output 1 and halt. Otherwise copy tape 2 to tape 1, erasing tape 2. By the promise, when $\text{PDISJ}_{N,p}^{\Pi} = 0$, we have $F_k(\mathbf{v}') = N' \leq N$; when $\text{PDISJ}_{N,p}^{\Pi} = 1$, we have

$$\begin{aligned} F_k(\mathbf{v}') &= N' - p + p^k \geq N' + 4\epsilon N - (4\epsilon N)^{1/k} \\ &\geq (1 + \epsilon)^2 N' + \epsilon N - (4\epsilon N)^{1/k} > (1 + \epsilon)^2 N' \end{aligned}$$

where the second inequality follows because $\epsilon \leq 1$ and the third follows because $\epsilon > 4^{1/(k-1)}/N$. The algorithm for $\text{PDISJ}_{N,p}^{\Pi}$ will output 1 if the value returned for F_k is greater than $(1 + \epsilon)N'$ and output 0 otherwise. \square

For our lower bound argument we will need the sequence of permutations Π to be of a special form. Let $p \geq 2$ and $N = mn$ where m and n are integers. A sequence $\Pi = (\pi_1, \dots, \pi_p)$ of permutations on $[N]$ has (*monotone*) *block-size* m if and only if there is a sequence $\Phi = (\phi_1, \dots, \phi_p)$ of permutations on $[m]$ such that $\pi_i(j) = (\phi_i(j') - 1)n + j''$ where $j = (j' - 1)n + j''$ with $j'' \in [n]$. That is each permutation π_i permutes blocks of length n in $[N]$ but leaves each block intact. In this case, we write $\text{PDISJ}_{n,p}^{m,\Phi}$ for $\text{PDISJ}_{N,p}^{\Pi}$.

Note that $\text{PDISJ}_{n,p}^{m,\Phi}$ can be viewed as the logical \vee of m independent copies of $\text{PDISJ}_{n,p}$ in which the input blocks for the different functions have been permuted by Φ .

3. Information flow for read/write streams

We first prove several results about the ‘‘information flow’’ in a read/write stream algorithm’s execution. These results

capture the information flow among the tapes in various stages of the computation. In this section, we only consider deterministic (r, s, t) -read/write stream algorithms.

Information flow in a read/write stream algorithm’s execution is captured via a *dependency graph* which shows the dependence of portions of the tapes on a given input string at various stages of a computation. This notion of dependency graphs is adapted from and simpler than the concept of skeletons that is used in [9, 8, 5]; although much simpler, it suffices for our purposes.

Recall that the input is a string $\mathbf{v} \in \{0, 1\}^*$ written on the first tape and that at any step only one of the t heads moves. The read/write stream algorithm A makes at most r reversals on input \mathbf{v} ; we can assume precisely r reversals without loss of generality. The dependency graph corresponding to \mathbf{v} , denoted by $\sigma(\mathbf{v})$, has $r + 2$ levels – level 0 corresponds to the beginning of the computation and level $r + 1$ corresponds the end of the computation. Level k for $1 \leq k \leq r$ encodes the dependency on the input of each of the t tapes immediately before the k -th reversal in the following manner: For $0 \leq k \leq r + 1$ there is one node at level k of $\sigma(\mathbf{v})$ for each tape cell that either contained a symbol of input \mathbf{v} or was visited at some time during the computation on input \mathbf{v} before the k -th reversal, or before the end of the computation for $k = r + 1$. The nodes at level k are ordered according to their positions on their corresponding cells on the tapes. Because of this we can view nodes of the dependency graph, interchangeably, as tape cells. There are pointers to each node at level k from each of the nodes in level $k - 1$ that it *depends* on.

The crucial observation made in [9] about read/write stream algorithms is the following: When a symbol is written in a particular cell by the read/write stream algorithm between its $k - 1$ -st and k -th reversal (i.e, at level k of $\sigma(\mathbf{v})$), what is being written in that cell can only depend on the current state and the t symbols currently being scanned by the read/write heads. However, the values of these t symbols were determined *before* the $k - 1$ -st reversal. This implies that any cell at level k depends either on t cells in level $k - 1$ (when it is overwritten in level k) or only on itself in level $k - 1$ (when it is intact in level k). The dependency graph thus consists of a layered directed graph of tape cells of in-degree either t or 1 representing the cell dependencies, where all the edges connect consecutive layers.

For some $b \geq 1$, suppose that \mathbf{v} is partitioned into b blocks of consecutive bits: $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_b)$. For every cell c in $\sigma(\mathbf{v})$, we write the *input dependency* of c as $\mathcal{I}_b(c) \subset \{1, \dots, b\}$ to denote the set of input blocks that it *depends* on (i.e, the set of i , for $1 \leq i \leq b$, such that there is a path from a cell in \mathbf{v}_i at level 0 to c). We note that the set $\mathcal{I}_b(c)$ depends on how we partition \mathbf{v} , which explains the subscript ‘ b ’ in the notation. Since in this paper we are only interested in those partitions into equal-length blocks,

this notation suffices; moreover, we will sometimes drop the subscript ‘ b ’ if it is clear from the context. It is immediate from the definition that for every cell c in level k for $0 \leq k \leq r + 1$, $|\mathcal{I}(c)| \leq t^k$.

For any cell c on a tape, we write $r(c)$ and $\ell(c)$ for the cells immediately to its right and to its left, respectively.

Proposition 3.1. *Suppose that an input \mathbf{v} is partitioned into b blocks: $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_b)$ for some $b \geq 1$. Let C be the set of cells on any one tape at any level k in $\sigma(\mathbf{v})$, for $0 \leq k \leq r + 1$. Fixing any $1 \leq i \leq b$, let $S = \{c \in C \mid i \in \mathcal{I}(c) \text{ and } i \notin \mathcal{I}(r(c))\}$ and $S' = \{c \in C \mid i \in \mathcal{I}(c) \text{ and } i \notin \mathcal{I}(\ell(c))\}$. Then $|S| = |S'|$ and $|S| \leq t^k$.*

Proof. The first part is obviously true, so we just prove the second part. We proceed by induction on k . For $k = 0$, this is true since the only cell in S is the one at the right boundary of the i -th input block.

For the induction step, consider any $k > 0$. At level k , a cell gets its input dependency from at most t cells on t tapes at level $k - 1$ that it depends on. Thus, for any cell $c \in S$, there must be some cell \tilde{c} from level $k - 1$ such that c depends on \tilde{c} , $i \in \mathcal{I}(\tilde{c})$, and either $r(c)$ depends on $r(\tilde{c})$ and $i \notin \mathcal{I}(r(\tilde{c}))$, or $r(c)$ depends on $\ell(\tilde{c})$ and $i \notin \mathcal{I}(\ell(\tilde{c}))$ (depending on whether c 's tape head direction and \tilde{c} 's tape head direction are the same or not between the $(k - 1)$ -st and k -th reversals). By induction, this happens at most $t \cdot t^{k-1} = t^k$ times. \square

For a set T , we write \mathcal{S}_T to denote the set of strings of length $|T|$ that are permutations of T . A string s of length $|s|$ is said to be the *interleaving* of another set of strings $\{s_1, \dots, s_t\}$ if there is a partition of $\{1, \dots, |s|\}$ into t subsets $\{q_1, \dots, q_t\}$ such that for every $1 \leq i \leq t$, $s_{|q_i} = s_i$, where $s_{|q_i}$ denotes the string obtained from s projected on coordinates in q_i only, and for every $j \in q_i$, the j -th entry of s is said to be *assigned* to s_i .

Consider the dependency graph $\sigma(\mathbf{v})$ associated with an input \mathbf{v} . For any level k in $\sigma(\mathbf{v})$ consider the sequence of nodes in $\sigma(\mathbf{v})$ corresponding to one of the tapes. An *input dependency string* C of this tape is any string in $\mathcal{S}_{\mathcal{I}(c_1)} \cdots \mathcal{S}_{\mathcal{I}(c_L)}$ where the cells of the tape are c_1, \dots, c_L in order. For any cell c_i , a string in $\mathcal{S}_{\mathcal{I}(c_i)}$ that is a substring of C is called a *cell portion* of C associated with c_i .

Proposition 3.2. *Let C be an input dependency string of any one tape at any level k in $\sigma(\mathbf{v})$, for $0 \leq k \leq r + 1$. Then C can be written as the interleaving of at most t^k monotone sequences so that for every such sequence s , there is at most one entry in every cell portion of C that is assigned to s .*

Proof. The general idea is taken from [8]. We proceed by induction on k . We will prove a somewhat stronger inductive claim, namely that the property above also holds for more general strings, namely those strings C in

$\bigcup_{a_i \geq 1} \mathcal{S}_{\mathcal{I}(c_1)}^{a_1} \cdots \mathcal{S}_{\mathcal{I}(c_L)}^{a_L}$, which we call the set of *extended dependency sequences* for a tape with cells c_1, \dots, c_L . Note that each of these strings can be partitioned into $\sum_{i=1}^L a_i$ cell portions, each of which corresponds to a string in $\mathcal{S}_{\mathcal{I}(c_i)}$ for some $1 \leq i \leq L$.

For $k = 0$, the only non-empty tape is the input tape and C itself is a monotone sequence. Thus this is true for $k = 0$.

For the induction step, suppose that the tape we are considering is the j -th tape, where $1 \leq j \leq t$. At level k the tape head visits consecutive cells in the j -th tape and the remaining cells are kept intact. Thus C can be written as $C = C' D C''$, where C' and C'' correspond to those cells that are intact from level $k - 1$ and D corresponds to those cells that are visited. For each of those former cells, its input dependency is unchanged from level $k - 1$, and for each of those latter cells, its input dependency is the union of those of the t cells it depends on. Thus D can be written as the interleaving of t sequences D_1, \dots, D_t , where sequence D_i for $1 \leq i \leq t$ denotes a substring of an extended input dependency of tape i from level $k - 1$. One observation here is that $C' D_j C''$ is also a substring of an extended input dependency string of tape j at level $k - 1$. By induction, each of D_1, \dots, D_t and $C' D_j C''$ can be written as the interleaving of at most t^{k-1} monotone sequences satisfying the requirement. Hence C can be written as the interleaving of at most t^k monotone sequences satisfying the requirement. \square

Proposition 3.3. *Suppose that an input \mathbf{v} is partitioned into b blocks $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_b)$ for some $b \geq 1$. For any cell c (at any level) in $\sigma(\mathbf{v})$, let $H(c) = \{\{i, j\} \mid 1 \leq i \neq j \leq b \text{ and } i, j \in \mathcal{I}(c)\}$. Then $|\bigcup_{c \in \sigma(\mathbf{v})} H(c)| \leq t^{3r+4}b$.*

Proof. Note that for any two input blocks \mathbf{v}_i and \mathbf{v}_j , if $i, j \in \mathcal{I}(c)$ for some cell c at any level $l < r + 1$, then $i, j \in \mathcal{I}(c')$ for some cell c' at level $r + 1$. Thus it suffices to consider the last level.

Fix any tape j for $1 \leq j \leq t$. Let C be an input dependency string of tape j at level $r + 1$. From Proposition 3.2, C can be decomposed into a set S of t^{r+1} interleaved monotone sequences as described there. For any cell c on this tape and for any sequence $s \in S$, we say that c stands at some *stage* i in s if i is the rightmost entry before or in the cell portion of c in C that is assigned to s . Define a table \mathcal{T} whose t^{r+1} rows correspond to the sequences in S and whose columns correspond to the cells in tape j as follows: \mathcal{T} has a number of columns equal to the number of cells with nonempty input dependency, where the columns are placed in the same left-to-right order as their corresponding cells. Each entry in \mathcal{T} records the stage at which its corresponding cell stands in its corresponding sequence. For each column \mathbf{c} corresponding to a cell c , let

$$H'(\mathbf{c}) = \{\{\mathcal{T}_{\mathbf{r}, \mathbf{c}}, \mathcal{T}_{\mathbf{r}', \mathbf{c}}\} \mid \mathbf{r} \neq \mathbf{r}' \text{ and } \mathcal{T}_{\mathbf{r}, \mathbf{c}} \neq \mathcal{T}_{\mathbf{r}', \mathbf{c}}\}.$$

Obviously $H(c) \subset H'(\mathbf{c})$ and $|H'(\mathbf{c})| \leq (t^{r+1})^2 = t^{2r+2}$.

For any two adjacent columns \mathbf{c} and \mathbf{c}' corresponding to two adjacent cells c and c' , respectively, if $H'(\mathbf{c}) \neq H'(\mathbf{c}')$, then there must be a sequence $s \in S$ such that c and c' stand at different stages in s . Since s is monotone, this happens at most b times. Since there are at most t^{r+1} sequences, there are at most bt^{r+1} different $H'(\mathbf{c})$ over all columns. Thus $|\cup_{c \in \sigma(\mathbf{v})} H(c)| \leq t \cdot t^{2r+2} \cdot bt^{r+1} = bt^{3r+4}$. \square

Proposition 3.4. *Suppose that an input \mathbf{v} is partitioned into $b = p \cdot m$ blocks: $\mathbf{v} = (\mathbf{v}_1, \dots, \mathbf{v}_{pm})$, for some $p \geq 2, m \geq 1$. For any sequence of permutations $\phi_1, \phi_2, \dots, \phi_p$ on $[m]$, there exists a set $I \subset [m]$ with $|I| \geq m - p \cdot t^{5r+8} \text{relsorted}(\{\phi_1, \dots, \phi_p\})$ satisfying the following property: for every $i \in I$, let $J_i = \{\phi_1(i), m + \phi_2(i), \dots, (p-1)m + \phi_p(i)\}$, then there is no cell c in $\sigma(\mathbf{v})$ such that $|J_i \cap \mathcal{I}(c)| > 1$.*

Proof. This is a generalization of an argument in [8] which gave a proof for the special case $p = 2$. As in the previous proposition, it suffices that we consider level $r + 1$ only. Let $Q = \{i \in [m] \mid \exists c : |J_i \cap \mathcal{I}(c)| > 1\}$. Then $I = [m] \setminus Q$. Thus we need to prove $|Q| \leq p \cdot t^{5r+8} \text{relsorted}(\{\phi_1, \dots, \phi_p\})$.

We partition Q into disjoint subsets such that $Q = \cup_{1 \leq p_1 < p_2 \leq p} Q_{p_1, p_2}$ and $Q_{p_1, p_2} \subseteq \{i \in Q \mid \exists c : (p_1 - 1)m + \phi_{p_1}(i), (p_2 - 1)m + \phi_{p_2}(i) \in \mathcal{I}(c)\}$ for any $1 \leq p_1 < p_2 \leq p$. By Proposition 3.3, there are at most pt^{3r+4} nonempty Q_{p_1, p_2} . Therefore there must be some $p_1 < p_2 \in [p]$ such that $|Q_{p_1, p_2}| \geq |Q| / (t^{3r+4} p)$. Fix these p_1 and p_2 .

Let $Q_{p_1, p_2} = \{i_1, \dots, i_q\}$, where $q = |Q_{p_1, p_2}|$. Let $C \in \{1, \dots, pm\}^*$ obtained by concatenating the input dependency of all t tapes, so that for any cell c , if $\mathcal{I}(c)$ contains both $(p_1 - 1)m + \phi_{p_1}(i)$ and $(p_2 - 1)m + \phi_{p_2}(i)$ for some $i \in Q_{p_1, p_2}$, then both of them are placed consecutively and in this order. By Proposition 3.2, C can be decomposed into a set S of $t \cdot t^{r+1} = t^{r+2}$ monotone sequences. Let π be a permutation on $\{1, \dots, q\}$ so that

$$(p_1 - 1)m + \phi_{p_1}(i_{\pi(1)}), (p_2 - 1)m + \phi_{p_2}(i_{\pi(1)}), \dots, \\ \dots, (p_1 - 1)m + \phi_{p_1}(i_{\pi(q)}), (p_2 - 1)m + \phi_{p_2}(i_{\pi(q)})$$

occur in the same order in C . Since there are q entries in input block p_1 and each of them must be in at least one sequence in S , there is some sequence $s \in S$ such that there are at least $q/|S| = q/t^{r+2}$ such entries in s . In other words, there exists a set $Q_1 \subset \{1, \dots, q\}$ of size at least q/t^{r+2} such that for every $j \in Q_1$, $(p_1 - 1)m + \phi_{p_1}(i_{\pi(j)}) \in s$, and since s is monotone, for every $j_1 < j_2 \in Q_1$, either $\phi_{p_1}(i_{\pi(j_1)}) < \phi_{p_1}(i_{\pi(j_2)})$ or $\phi_{p_1}(i_{\pi(j_1)}) > \phi_{p_1}(i_{\pi(j_2)})$, depending on the monotony of s . Let the indices in Q_1 be $j_1 < \dots < j_{q_1}$. Consider the following list of entries

$$(p_2 - 1)m + \phi_{p_2}(i_{\pi(j_1)}), \dots, (p_2 - 1)m + \phi_{p_2}(i_{\pi(j_{q_1})}),$$

which occurs in this order in C . As before, there must be at least one sequence $s' \in S$ such that there are at least

$q_1/|S| = q_1/t^{r+2}$ such entries in s' . In other words, there exists a set $Q_2 \subset Q_1$ of size at least $q_1/t^{r+2} = q/t^{2r+4}$ such that for every $l \in Q_2$, $(p_2 - 1)m + \phi_{p_2}(i_{\pi(l)}) \in s'$, and hence since s' is monotone, for every $l_1 < l_2 \in Q_2$, either $\phi_{p_2}(i_{\pi(l_1)}) < \phi_{p_2}(i_{\pi(l_2)})$ or $\phi_{p_2}(i_{\pi(l_1)}) > \phi_{p_2}(i_{\pi(l_2)})$, depending on the monotony of s' . Therefore, by definition of sortedness we see that

$$\text{relsorted}(\{\phi_1, \dots, \phi_p\}) \geq \text{sortedness}(\phi_{p_1}, \phi_{p_2}) \geq q/t^{2r+4}$$

which gives $q \leq \text{relsorted}(\{\phi_1, \dots, \phi_p\})t^{2r+4}$ and hence concludes the proposition. \square

4. Simulation of read/write stream algorithms by communication protocols

Let $p \geq 2, m \geq 1$ and $\Phi = (\phi_1, \dots, \phi_p)$ be a sequence of p permutations defined on $[m]$. Let X be a non-empty set and $Y = (Y_1, \dots, Y_p) \in X^p$.

For each $i \in [m]$ and $\rho \in X^{(m-1)p}$, we define $J_i = \{\phi_1(i), m + \phi_2(i), \dots, (p-1)m + \phi_p(i)\}$ and $\mathbf{v} = \mathbf{v}(Y, i, \rho, \Phi) = (\mathbf{v}_1, \dots, \mathbf{v}_{pm}) \in X^{pm}$ such that $\mathbf{v}_{(j-1)m + \phi_j(i)} = Y_j$ for every $j \in [p]$, and $\mathbf{v}_{\{1, \dots, pm\} - J_i} = \rho$. Let A be a deterministic read/write stream algorithm defined on X^{pm} . Let $\sigma(\mathbf{v})$ be the dependency graph induced by A on \mathbf{v} and $I_{\mathbf{v}} \subseteq [m]$ be the set of input indexes defined by $\sigma(\mathbf{v})$ as described in Proposition 3.4.

The following theorem will be used to show that by simulating an efficient deterministic read/write stream algorithm for f_{Φ}^{\vee} , one can derive efficient p -party number-in-hand communication protocols for a variety of embeddings of f in f_{Φ}^{\vee} .

Theorem 4.1. *Let $p \geq 2, m \geq 1$ and $\Phi = (\phi_1, \dots, \phi_p)$ be a sequence of permutations on $[m]$. Let $X \neq \emptyset$ with $n = \lceil \log_2(X) \rceil$. Suppose that A is a deterministic (r, s, t) -algorithm defined on X^{pm} . Then there is a constant $c > 0$ depending on t such that, for each $i \in [m]$ and $\rho \in X^{p(m-1)}$ there is a p -party number-in-hand protocol $P_{i, \rho}$ in which player j has access to $Y_j \in X$ (and implicitly i and ρ) for each j , each $P_{i, \rho}$ communicates at most $O(t^{cr} p(s + \log pmn))$ bits, and for each i, ρ and $\mathbf{v} = \mathbf{v}(Y, i, \rho, \Phi)$ there is a set $I'_{i, \rho, \mathbf{v}} \subseteq [m]$ containing $I_{\mathbf{v}}$ such that if $i \notin I'_{i, \rho, \mathbf{v}}$ then $P_{i, \rho}$ on input \mathbf{v} outputs “fail” and if $i \in I'_{i, \rho, \mathbf{v}}$ then $P_{i, \rho}$ on input \mathbf{v} outputs the value $A(\mathbf{v})$.*

Before going to the proof we need the following lemma.

Lemma 4.2. *When A terminates, the total length of all tapes used by A is at most $2^{(r+1)s} pmn$.*

Proof. The initial total length is clearly pmn . It is also clear that immediately after each reversal, the total length is multiplied by at most 2^s . The lemma follows. \square

Proof of Theorem 4.1. We describe and then analyze the protocol. Each player first constructs $\mathbf{v} = \mathbf{v}(Y, i, \rho, \Phi)$ and

then executes A on \mathbf{v} . Note that all players can access the whole input \mathbf{v} except for the p blocks holding Y_1, \dots, Y_p , each of which is known to exactly one player. Since no player knows the whole input, in order to correctly simulate A , they need to communicate during the simulation. Along the way, each player gradually constructs and keeps a copy of $\sigma(\mathbf{v})$. Each keeps track of the level (number of reversals) in $\sigma(\mathbf{v})$ that A is currently working on and the machine state of A . Essentially, for every tape cell at every level in $\sigma(\mathbf{v})$ written by A , the players record whether (1) the contents of the cell can be computed by everyone, or (2) the contents of the cell can only be computed by a specific player.

Those cells of type (1) are those cells c such that $(j - 1)m + \phi_j(i) \notin \mathcal{I}_{pm}(c)$ for any $1 \leq j \leq p$. For each of these cells, each of the players records: the machine state immediately before overwriting the cell, and the (at most t) cells of the previous level on which this cell depends. Note that those cells that a cell of type (1) depends on are also type (1) cells. It is clear that by recursion, every player can compute the contents of each of these cells as needed.

Those cells of type (2) are those that depend on some input held by a particular player. Consider a cell c such that $(j - 1)m + \phi_j(i) \in \mathcal{I}_{pm}(c)$ for some $j \in [p]$. Each player records that this cell depends on player j . We will show later what information player j needs to record so that she can compute the contents of c herself.

Note that there is another type of cell, whose contents depend on the inputs from more than one player. As soon as the simulation discovers one of these cells, it will stop and the protocol outputs “fail”. We will explain more about this point later.

The simulation proceeds as follows. Each player executes A step by step. At every new step in which all the t tape heads are to read cells of type (1) only, every player can compute the contents of the t cells without any communication. Since each of them holds the current machine state, they can compute which one of the t tapes is written and the moves and the content of the write. Each of them thus records, for the overwritten cell, that it is of type (1) as well as the tape heads and the machine state. To end this step, each of the players also has the new machine state.

The more interesting case is when at a new step, at least one of the tape heads is to read at least one cell of type (2) and all of the type (2) cells depend on a player j . All players will then wait for player j to communicate. Player j will proceed as follows. As long as at least one of the tape heads still reads a cell depending on her or the algorithm does not make any reversal, she proceeds with the simulation, and clearly has sufficient information to do so. Along the way, for every cell she overwrites, she records the machine state and all the tape head positions for that cell, so that she can compute the cell later when needed. This process stops when the algorithm comes to a new step in which

either all the tape heads are to read a cell of type (1), or at least one of the tape heads depends on another player, or one of the tape heads reverses its direction. When this process stops, player j broadcasts: (a) all t updated tape head positions and directions, and (b) the new machine state. Since there has been no reversal, all other players know precisely which cells were visited by player j and they mark all those overwritten cells, which are all of the same level in $\sigma(\mathbf{v})$, as of type (2) and depending on j . Therefore, all players now have sufficient information to proceed.

The last case is when at a new step, at least two of the tape heads are to read cells of type (2) and these two cells depend on two different players. In this case, all players stop the simulation and output “fail”. It is clear that if $i \in I_{\mathbf{v}}$, by Proposition 3.4, this case will never happen.

When A terminates, the protocol will output exactly as A does. It remains to compute the communication cost.

We need to bound the cost of each communication and the number of times a communication occurs. From Lemma 4.2, the cost of one communication is $t \log(2^{rs}pn) + s = O(rs + \log pmn)$, where the hidden constant depends on t . When one player communicates, one of the tape heads has either just reversed or just moved from a cell depending on her to another cell that does not. The former means that the algorithm comes to the next level, which happens at most r times. By Proposition 3.1 (with $b = p$), the latter occurs at most $t^{r+1} + t^r + \dots + 1$ times for a single tape and single player. Summing over all tapes and all p players, this occurs at most pt^{r+3} times in total. \square

5. Bounds for disjointness and frequency moments

Using Lemma 2.3 and the fact that $\text{PDISJ}_{n,p}^{m,\Phi}$ is a special case of $\text{PDISJ}_{N,p}^{\Pi}$ where $N = mn$, we will lower bound the frequency moment problems by giving lower bounds for the $\text{PDISJ}_{n,p}^{m,\Phi}$ problem. This problem was previously studied in [5] for the special case $p = 2$; our bounds extend those in [5] for any $p \geq 2$ and also improve the bounds for $p = 2$.

Lemma 5.1. *There is a positive constant c such that given a randomized (r, s, t) -read/write stream algorithm A for $\text{PDISJ}_{n,p}^{m,\Phi}$ on X^{pm} , where $X = \{0, 1\}^n$ and $\frac{pt^{5r+8} \text{relsorted}(\Phi)}{m} = d < 1$, with error at most δ , there is a randomized public-coin p -party number-in-hand protocol P for $\text{PDISJ}_{n,p}$ of communication complexity $O(t^{cr}p(s + \log pmn))$ with error at most $\delta + d(1 - \delta)$.*

Proof. Suppose that A uses at most Γ random bits in its execution. For any string $\mathfrak{R} \in \{0, 1\}^{\Gamma}$, let $A_{\mathfrak{R}}$ denote the deterministic algorithm obtained from A using \mathfrak{R} as source of randomness. Following the view of X as representing the power set of $[n]$, we use $|x|$ to denote the number of 1’s in $x \in X$ and $\binom{[n]}{\ell}$ to denote the set of $x \in X$ with $|x| = \ell$.

We consider the following public-coin randomized communication protocol \mathcal{P} . On input $Y = \{Y_1, \dots, Y_p\} \in X^p$:

1. Each player j broadcasts $|Y_j|$.
2. The players use the public coins to
 - (a) uniformly and randomly generate $\rho \in (f^{-1}(0) \cap [\times_{j=1}^p \binom{[n]}{|Y_j|}])^{m-1} \subseteq X^{p(m-1)}$,
 - (b) choose $1 \leq i \leq m$ uniformly and randomly,
 - (c) uniformly and randomly generate $\mathfrak{R} \in \{0, 1\}^\Gamma$,
 - (d) uniformly and randomly generate a permutation $\phi : [n] \rightarrow [n]$.
3. Each player j computes $Y'_j = \phi(Y_j)$.
4. The players run protocol $P_{i,\rho}$ based on $A_{\mathfrak{R}}$ on inputs $Z = Z(Y'_1, \dots, Y'_p, i, \rho, \Phi)$, as described in Theorem 4.1
5. If $P_{i,\rho}$ outputs “fail” then output 1; else output what $P_{i,\rho}$ does.

We analyze this protocol \mathcal{P} . Let $\sigma(\mathbf{v})$ be the dependency graph induced by $A_{\mathfrak{R}}$ on input Z and $I_Z \subseteq [m]$ be the set of input indexes defined by $\sigma(\mathbf{v})$ as described in Proposition 3.4. Let $I' = I'_{i,\rho,Z}$ be the set of input positions induced by $A_{\mathfrak{R}}$ given the choice of i, ρ and input Z as described in Theorem 4.1. The correctness of the protocol depends on whether $i \in I'$ and whether $A_{\mathfrak{R}}(Z)$ is correct.

First we consider the case that the sets Y_1, \dots, Y_p are not disjoint; i.e., $\text{PDISJ}_{n,p}(Y) = 1$. In this case if $P_{i,\rho}$ outputs “fail”, the protocol always outputs correctly. Otherwise it will output what A does. In other words, we have

$$\begin{aligned} & \Pr[\mathcal{P}(Y) = 1 \mid \text{PDISJ}_{n,p}(Y) = 1] \\ &= \Pr[\mathcal{P}(Y) = 1 \text{ and } i \notin I' \mid \text{PDISJ}_{n,p}(Y) = 1] \\ & \quad + \Pr[\mathcal{P}(Y) = 1 \text{ and } i \in I' \mid \text{PDISJ}_{n,p}(Y) = 1] \\ &= \Pr[i \notin I' \mid \text{PDISJ}_{n,p}(Y) = 1] \\ & \quad + \Pr[A_{\mathfrak{R}}(Z) = 1 \text{ and } i \in I' \mid \text{PDISJ}_{n,p}(Y) = 1] \\ &\geq \Pr[A_{\mathfrak{R}}(Z) = 1 \mid \text{PDISJ}_{n,p}(Y) = 1] \geq 1 - \delta. \end{aligned}$$

Next we consider the case the sets Y_1, \dots, Y_p are disjoint. After being re-mapped by ϕ , the sets Y'_1, \dots, Y'_p are also disjoint and uniformly distributed in $f^{-1}(0) \cap [\times_{j=1}^p \binom{[n]}{|Y_j|}]$. Hence Z is uniformly distributed over $(f^{-1}(0) \cap [\times_{j=1}^p \binom{[n]}{|Y_j|}])^m$. Therefore I_Z is statistically independent from i . By Theorem 4.1, $I_Z \subseteq I'$ and by Proposition 3.4, $|I_Z|$ is always at least $(1-d)m$. Thus the probability that $i \in I'$ is at least $1-d$. Formally, we have

$$\begin{aligned} & \Pr[\mathcal{P}(Y) = 0 \mid \text{PDISJ}_{n,p}(Y) = 0] \\ &\geq \Pr[A_{\mathfrak{R}}(Z) = 0 \mid \text{PDISJ}_{n,p}(Y) = 0] \\ & \quad \cdot \Pr[\mathcal{P}(Y) = 0 \mid A_{\mathfrak{R}}(Z) = 0, \text{PDISJ}_{n,p}(Y) = 0] \\ &\geq (1-\delta) \Pr[\mathcal{P}(Y) = 0 \mid A_{\mathfrak{R}}(Z) = 0, \text{PDISJ}_{n,p}(Y) = 0] \\ &= (1-\delta) \Pr[i \in I' \mid A_{\mathfrak{R}}(Z) = 0, \text{PDISJ}_{n,p}(Y) = 0] \\ &\geq (1-\delta) \Pr[i \in I_Z \mid A_{\mathfrak{R}}(Z) = 0, \text{PDISJ}_{n,p}(Y) = 0] \\ &= (1-\delta) \Pr[i \in I_Z \mid \text{PDISJ}_{n,p}(Y) = 0] \geq (1-\delta)(1-d). \end{aligned}$$

The communication complexity is $p \log_2 n$ plus the complexity of $P_{i,\rho}$ which completes the lemma. \square

Lemma 5.2. *Let $\delta < 1/4$, $t \geq 1$, $1/2 > \beta > 0$, $\epsilon > \alpha > 0$, and $m, n, p > 1$ so that $p \leq m^{\frac{1}{2}-\alpha} = n^\beta$. Let $N = mn$. Then for large enough N , there is a constant $a > 0$ with the following property: if $r \leq a \log N$ and $s = o(N^{1-\frac{4\beta}{2\beta+1}-\epsilon})$, then there is no randomized (r, s, t) -read/write stream algorithm with error at most δ for $\text{PDISJ}_{n,p}^{m,\Phi^*}$ on $\{0, 1\}^{pmn}$.*

Proof. Suppose by contradiction that there is such an algorithm A. By Lemma 5.1, there is a public-coin randomized number-in-hand communication protocol for $\text{PDISJ}_{n,p}$ with complexity $O(t^{cr}ps)$ for some positive constant c with error at most $d + (1-d)\delta$, where $d = \frac{pt^{5r+8} \text{reلسorted}(\Phi^*)}{m}$. Since $r \leq a \log N = a(1 + \frac{1-2\alpha}{2\beta}) \log m$ and $\text{reلسorted}(\Phi^*) \leq 2e\sqrt{m}$, we have $d + (1-d)\delta \leq 2\delta < 1/2$ for a sufficiently small depending on δ, t , and α .

By [7], this communication complexity must be $\Omega(\frac{n}{p \log p})$. This gives us, for some constant ϵ'' depending on a and α that $s = \Omega(n^{1-2\beta-\epsilon''}) = \Omega(N^{\frac{(1-2\beta-\epsilon'')(1-2\alpha)}{1+2\beta-2\alpha}}) = \Omega(N^{1-\frac{4\beta}{2\beta+1}-\epsilon})$. \square

This immediately implies a lower bound on $\text{PDISJ}_{N,p}^{\Pi^*}$ where Π^* is the extension of Φ^* to a sequence of p permutations on $[N]$.

Theorem 5.3. *Let $\delta < 1/4$, $t \geq 1$, $\epsilon > 0$, and $N, p > 1$ so that $p \leq N^\gamma$ for $\gamma < \frac{1}{4}$. Then for large enough N , there exists sequence Π^* of p permutations on $[N]$ such that there is no randomized $(o(\log N), N^{1-4\gamma-\epsilon}, t)$ -read/write stream algorithm with error at most δ for $\text{PDISJ}_{N,p}^{\Pi^*}$.*

Proof. Follows from Lemma 5.2 where we set $N = mn$, $\Pi^* = (\pi_1^*, \dots, \pi_p^*)$ to be the extension of Φ^* to p permutations on $[N]$ as described in Section 2, and set α, β so that $\frac{1}{\gamma} = \frac{1}{\beta} + \frac{2}{1-2\alpha}$ and α is sufficiently small. \square

By Theorem 5.3 and Lemma 2.3(a) we obtain the following lower bounds for approximating F_k and F_∞^* .

Corollary 5.4. *Let $k > 1$, $t \geq 1$, $\eta > 0$, and $1 \geq \epsilon \geq 1/N$. Then there is no randomized $(o(\log N), O(\frac{1}{\epsilon^{4/k}} N^{1-\frac{4}{k}-\eta}), t)$ -read/write stream algorithm that with probability at least $3/4$ outputs an approximation of F_k within a factor of $(1+\epsilon)$ on an input stream of up to N elements from $[N]$.*

Proof. Let ζ satisfy $\epsilon = N^{-\zeta}/4$. Then setting $p = \lceil N^{(1+\zeta)/k} \rceil = \lceil (4\epsilon N)^{1/k} \rceil$ and applying Lemma 2.3(a) and Theorem 5.3 we obtain that no randomized read/write stream algorithm with $o(\log N)$ reversals, t tapes, and space $O(N^{1-4(1-\zeta)/k-\eta})$ can compute F_k within a $(1+\epsilon)$ factor. Replacing $N^{-\zeta}$ by 4ϵ yields the claimed bound. \square

Letting $\epsilon = 1$ in the above implies that there is no factor 2 approximation to F_k for $k > 4$ that uses small space and a small number of reversals in the read/write streams model.

Corollary 5.5. *Let $k > 4$, $t \geq 1$, and $\eta > 0$. Then there is no randomized $(o(\log N), O(N^{1-\frac{4}{k}-\eta}), t)$ -read/write stream algorithm that with probability at least $3/4$ outputs an approximation of F_k within a factor of 2 on an input stream of up to N elements from $[N]$.*

By similar means we can derive improved lower bounds for computing F_∞^* .

Corollary 5.6. *Let $t \geq 1$, and $\eta > 0$. Then there is no randomized $(o(\log N), O(N^{1-\eta}), t)$ -read/write stream algorithm that with probability at least $3/4$ outputs an approximation of F_∞^* within a factor of 2 on an input stream of up to N elements from $[N]$.*

We also derive lower bounds for the case that $k < 1$ using Theorem 5.3 and Lemma 2.3(b).

Corollary 5.7. *Let $k < 1$, $t \geq 1$, $\eta > 0$, and $1/N \leq \epsilon < 1/N^{3/4+\eta}$. Then there is no randomized $(o(\log N), O(\frac{1}{\epsilon^4 N^{3+\eta}}), t)$ -read/write stream algorithm that with probability at least $3/4$ outputs an approximation of F_k within a factor of $(1 + \epsilon)$ on an input stream of up to N elements from $[N]$.*

Proof. Define ζ so that $\epsilon = N^{-\zeta}/3$. Then for N sufficiently large as a function of k , if $p = N^{1-\zeta} = 3\epsilon N$ then $p - p^k \geq 2\epsilon N$ so Lemma 2.3(b) and Theorem 5.3 imply that no randomized read/write stream algorithm with $o(\log N)$ reversals, t tapes, and space $O(N^{1-4(1-\zeta)-\eta})$ can compute F_k within a $(1 + \epsilon)$ factor as described. Replacing $N^{-\zeta}$ by 3ϵ yields the claimed bound. \square

Our lower bound for $\text{PDISJ}_{N,p}^\Pi$ is only interesting when $N = \omega(p^4)$.² This is because in order for the reduction from $\text{PDISJ}_{n,p}^{m,\Phi}$ to work (Lemma 5.2), we need $N = nm$ and both $n = \omega(p^2)$ and $m = \omega(p^2)$. The condition $n = \omega(p^2)$ is induced by the communication complexity lower bound for $\text{PDISJ}_{n,p}$, which is optimal up to a logarithmic factor. The following lemma shows that a condition requiring m to be polynomially larger than p is also necessary and that the above technique cannot yield bounds for constant factor approximations for $k < 3.5$.

²Note that if N is $O(p^2)$ there is a simple deterministic algorithm for $\text{PDISJ}_{N,p}^\Pi$ for any Π . First check that the total size of the p subsets is at most N ; otherwise output 1. Then scan for a subset of size s at most $N/p = O(p)$ and look for these p elements in each of the other subsets in turn. Then split these s elements into $p-1$ groups of consecutive elements of constant size and look for the elements in the i -th group in the i -th of the other subsets. It is clear that the algorithm can be implemented using only two tapes, $O(\log n)$ space, and $O(1)$ reversals.

Lemma 5.8. *For integer $m < p^{3/2}/64$, any integer n , and for any $\Phi = (\phi_1, \dots, \phi_p)$ defined on $[m]$, there is a deterministic $(2, O(\log(mnp)), 2)$ -read/write stream algorithm computing $\text{PDISJ}_{n,p}^{m,\Phi}$.*

To produce the algorithm claimed in Lemma 5.8, we need to show the following property of permutations that does not appear to have been considered previously. Its proof is inspired by Seidenberg's proof of the well-known theorem of Erdős and Szekeres (cf. [15]) which shows that any pair of permutations must have relative sortedness at least \sqrt{m} . The difference is that with three permutations we can now ensure that the sequences appear in the same order in two of them rather than one possibly being reversed.

Lemma 5.9. *Let ϕ_1, ϕ_2 , and ϕ_3 be permutations on $[m]$. Then there is some pair $1 \leq a < b \leq 3$ such that $\phi_a(1), \dots, \phi_a(m)$ and $\phi_b(1), \dots, \phi_b(m)$ have a common subsequence of length at least $m^{1/3}$.*

Proof. Suppose by contradiction that there are no such a and b . For every $i \in [m]$, let $\ell_i \in [m]^{1/3}$, where $m' = \lceil m^{1/3} - 1 \rceil$, be defined as follows: $\ell_i[1]$ is the length of the longest common subsequence of $\phi_1(1), \dots, \phi_1(s)$ and $\phi_2(1), \dots, \phi_2(t)$, where $\phi_1(s) = \phi_2(t) = i$, and $\ell_i[2]$ and $\ell_i[3]$ are defined analogously for the other two pairs ϕ_2, ϕ_3 , and ϕ_1, ϕ_3 , respectively.

Now for any $i \neq j \in [m]$, we must have $\ell_i \neq \ell_j$. This is because there must be some pair, say ϕ_1 and ϕ_2 , such that either i occurs before j in both sequences or j occurs before i in both. In the first case $\ell_i[1] < \ell_j[1]$ and in the second case $\ell_i[1] > \ell_j[1]$.

However since $m' < m^{1/3}$, the number of different ℓ_i over all $i \in [m]$ is strictly $< m$ which is a contradiction. \square

It is not hard to show that the above lemma is tight, even for any four permutations.

Proof of Lemma 5.8. Given Φ there exist $L_1, L_2, \dots, L_{p/3}$ defined as follows: L_1 is a common subsequence of two of ϕ_1, ϕ_2 , and ϕ_3 , of length at least $m^{1/3}$ given by Lemma 5.9; L_2 is a common subsequence of two of ϕ_4, ϕ_5 , and ϕ_6 that is disjoint from L_1 and of length at least $(m - |L_1|)^{1/3}$; L_3 is a common subsequence of two of ϕ_7, ϕ_8 , and ϕ_9 disjoint from $L_1 \cup L_2$ and of length at least $(m - |L_1| - |L_2|)^{1/3}$, and so on, with no index $i \in [m]$ appearing in more than one sequence. For each of the L_j let a_j and b_j denote the indices of the two permutations having the common subsequence L_j . The number of elements that do not appear in L_1, \dots, L_ℓ is at most m_ℓ where m_ℓ is defined by the recurrence with $m_0 = m$ and $m_{j+1} = m_j - \lceil m_j^{1/3} \rceil$ for $j > 0$. If $m < p^{3/2}/64$, then $p \geq (64m)^{2/3} = 16m^{2/3}$. Now if $m_{p/8} > m/8$ then at least $(m/8)^{1/3} = m^{1/3}/2$ elements have been removed for each of $p/8 = 2m^{2/3}$ steps which implies $m_{p/8} = 0$, which is a contradiction. Repeating this argument reduces

m_j to at most $m/64$ after another $2(m/8)^{2/3} = m^{2/3}/2 = p/32$ steps. Repeating this eventually yields that $m_{p/3} = 0$, which implies that every $i \in [m]$ is in exactly one L sequence.

The algorithm copies the input to tape 2 leaving both heads at the right end of the tape. It will use the head on tape 1 to scan the blocks for the players and the head on tape 2 to scan the corresponding blocks for the even-numbered players. It will solve the disjointness problems for each block in the common subsequences $L_{p/3}, \dots, L_2, L_1$, in turn. If an intersection is found in some pair of corresponding blocks in these sequences then the output is 1; otherwise, the output is 0. The promise ensures that, for each of the m subproblems, to check for a given common element it suffices to compare the blocks for a single pair of players. Since every $i \in [m]$ appears in some L_j , if we can position the heads to check the corresponding blocks then we can compute each of the m disjointness subproblems exactly and hence $\text{PDISJ}_{n,p}^{m,\Phi}$.

The positioning of the tape heads to execute these comparisons can be achieved by suitably hardwiring the sequences $L_{p/3}, \dots, L_1$ into the state transition function of the algorithm; this requires only $O(\log pmn)$ bits of state as required. We leave the details to the full paper. \square

6. Open questions

The general question that we have begun to answer here is: for what natural approximation problems does the read/write streams model (with $o(\log N)$ passes) add significant power to the data streams model? We have mostly but not fully resolved the case of frequency moments – can one close the gap between the upper and lower bounds for computing $\text{PDISJ}_{N,p}^{\Pi}$ and approximating F_k ? As we have shown, we are not far from the limit on lower bounds using the blocked and permuted version of disjointness, $\text{PDISJ}_{n,p}^{m,\Phi}$; moreover, it is not clear how a simulation like that in Theorem 4.1 (which appears to be optimal up to a polylogarithmic factor) can be made to work for more general instances of $\text{PDISJ}_{N,p}^{\Pi}$.

Optimizing our upper bound for $\text{PDISJ}_{n,p}^{m,\Phi}$ raises the following interesting combinatorial question: Given a set of k permutations on $[m]$, what is the length of the longest common subsequence that can be guaranteed between some pair of these k permutations as a function of m and k ? We originally conjectured that subsequences of length $m^{1/2-o(1)}$ must exist for $k = O(\log m)$, which would have shown the approximate optimality of our algorithm, but recently $\tilde{O}(m^{1/3})$ upper bounds have been shown for this case [4].

Acknowledgements We would like to thank T.S. Jayram and Atri Rudra for useful discussions and an anonymous referee for helpful comments.

References

- [1] N. Alon, Y. Matias, and M. Szegedy. The space complexity of approximating the frequency moments. *J. Comput. Sys. Sci.*, 58(1):137–147, 1999.
- [2] B. Babcock, S. Babu, M. Datar, M. R., and J. Widom. Models and issues in data stream systems. In *Proc. 21st ACM PODS*, pages 1–16, 2002.
- [3] Z. Bar-Yossef, T. Jayram, R. Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *J. Comput. Sys. Sci.*, 68(4):702–732, 2004.
- [4] P. Beame, E. Blais, and D.-T. Huynh-Ngoc. Longest common subsequences in sets of permutations. In preparation.
- [5] P. Beame, T. S. Jayram, and A. Rudra. Lower bounds for randomized read/write stream algorithms. In *Proc. 39th ACM STOC*, pages 689–698, 2007.
- [6] L. Bhuvanagiri, S. Ganguly, D. Kesh, and C. Saha. Simpler algorithms for estimating frequency moments of data streams. In *Proc. Seventeenth ACM-SIAM SODA*, pages 708–713, 2006.
- [7] A. Chakrabarti, S. Khot, and X. Sun. Near-optimal lower bounds on the multi-party communication complexity of set disjointness. In *Eighteenth IEEE Conference on Computational Complexity*, pages 107–117, 2003.
- [8] M. Grohe, A. Hernich, and N. Schweikardt. Randomized computations on large data sets: Tight lower bounds. In *Proc. 25th ACM PODS*, pages 243–252, 2006.
- [9] M. Grohe and N. Schweikardt. Lower bounds for sorting with few random accesses to external memory. In *Proc. 24th ACM PODS*, pages 238–249, 2005.
- [10] P. Indyk and D. P. Woodruff. Tight lower bounds for the distinct elements problem. In *Proc. 44th IEEE FOCS*, pages 283–292, 2003.
- [11] P. Indyk and D. P. Woodruff. Optimal approximations of frequency moments of data streams. In *Proc. 37th ACM STOC*, pages 202–208, 2005.
- [12] S. Muthukrishnan. Data streams: Algorithms and applications. *Foundations and Trends in Theoretical Computer Science*, 1(2), 2006.
- [13] A. A. Razborov. On the distributional complexity of disjointness. *Theoretical Computer Science*, 106(2):385–390, 1992.
- [14] M. E. Saks and X. Sun. Space lower bounds for distance approximation in the data stream model. In *Proc. 34th ACM STOC*, pages 360–369, 2002.
- [15] J. M. Steele. Variations on the monotone subsequence problem of Erdős and Szekeres. In Aldous, Diaconis, and Steele, editors, *Discrete Probability and Algorithms*, pages 111–132. Springer-Verlag, 1995.
- [16] D. P. Woodruff. Optimal space lower bounds for all frequency moments. In *Proc. Fifteenth ACM-SIAM SODA*, pages 167–175, 2004.