

Data Structure Lower Bounds for Document Indexing Problems*

Peyman Afshani¹ and Jesper Sindahl Nielsen²

1 MADALGO, Department of Computer Science, Aarhus University, Denmark[†]
peyman@madalgo.au.dk

2 MADALGO, Department of Computer Science, Aarhus University, Denmark[‡]
jasn@cs.au.dk

Abstract

We study data structure problems related to document indexing and pattern matching queries and our main contribution is to show that the pointer machine model of computation can be extremely useful in proving high and unconditional lower bounds that cannot be obtained in any other known model of computation with the current techniques. Often our lower bounds match the known space-query time trade-off curve and in fact for all the problems considered, there is a very good and reasonable match between our lower bounds and the known upper bounds, at least for some choice of input parameters.

The problems that we consider are *set intersection* queries (both the *reporting* variant and the *semi-group counting* variant), indexing a set of documents for *two-pattern* queries, or *forbidden-pattern* queries, or *queries with wild-cards*, and indexing an input set of *gapped-patterns* (or *two-patterns*) to find those matching a document given at the query time.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases Data Structure Lower Bounds, Pointer Machine, Set Intersection, Pattern Matching

Digital Object Identifier 10.4230/LIPIcs.ICALP.2016.93

1 Introduction

In this paper we study a number of data structure problems related to document indexing and prove space and query time lower bounds that in most cases are almost tight. Unlike many of the previous lower bounds, we disallow random accesses by working in the pointer machine model of computation, however, we obtain high and unconditional space and query time lower bounds that almost match the best known data structures for all the problems considered; at the moment, obtaining such unconditional and tight lower bounds in other models of computation is a hopelessly difficult problem. Furthermore, compared to the previous lower bounds in the area, our lower bounds probe deeper and thus are much more informative. Consequently, these results show the usefulness of the pointer machine model.

Document indexing is an important problem in the field of information retrieval. Generally, the input is a collection of documents $\mathcal{D} = \{d_1, d_2, \dots, d_D\}$ with a total length of n characters

* A full version of this paper is available, see [5].

[†] Research funded by MADALGO, Center for Massive Data Algorithmics, a Center of the Danish National Research Foundation, grant DNRF84

[‡] Research funded by MADALGO, Center for Massive Data Algorithmics, a Center of the Danish National Research Foundation, grant DNRF84



© Peyman Afshani and Jesper Sindahl Nielsen;
licensed under Creative Commons License CC-BY

43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016).

Editors: Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi;
Article No. 93; pp. 93:1–93:15



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Table 1** At every row, the 3rd cell presents our space lower bound for data structures that have a query time bounded by the 2nd cell. PM stands for the pointer machine model. n is the input size and t is the output size.

Problem	Query Bound	Space Lower Bound	Assumptions
2P, FP, SI, 2FP (counting)	$Q(n)$	$\Omega(n^{2-o(1)}/Q^2(n))$	Semi-group
2P, FP, SI, 2FP (reporting)	$Q(n) + \mathcal{O}(t)$	$\Omega(n^{2-o(1)}/Q(n))$	PM
2P, FP, SI, 2FP (reporting)	$\mathcal{O}((nt)^{\frac{1}{2}-\alpha} + t)$	$\Omega\left(n^{\frac{1+6\alpha}{1+2\alpha}-o(1)}\right)$	PM, $\alpha > 0$ a parameter
WCI (reporting)	$Q(n, \kappa) + \mathcal{O}(t)$	$\Omega\left(\frac{n}{\kappa} \Theta\left(\frac{\log n}{\kappa \log Q(n, \kappa)}\right)^{\kappa-1}\right)$	PM, κ wild-cards, $\kappa \leq \log_{Q(n, \kappa)} n$
WCI (reporting)	$\mathcal{O}(2^{\kappa/2} + t)$	$\Omega\left(n^{1+\Theta(1/\log k)}\right)$	PM, κ wild-cards, $3\sqrt{\log n} \leq \kappa = o\left(\frac{\log n}{\log \log n}\right)$
κ -GPI	$o\left(\frac{D\gamma^\kappa}{(2+2\log D)^{\kappa+1}}\right)$	$n^{\Omega(\log^{1/(2\kappa)} n)}$	PM, $\alpha_i = 0, \beta_i = \gamma$ $\kappa = o\left(\frac{\log \log n}{\log \log \log n}\right)$

and usually the goal is to index them such that given a query, all the documents matching the query can be either found efficiently (the *reporting variant*) or counted efficiently (the *searching variant*). When the query is just one text pattern, the problem is classical and well-studied and there are linear space solutions with optimal query time [33]. Not surprisingly, there have been various natural extensions of this problem. We summarize the problems we study and our results below.

1.1 New and Previous Results

1.1.1 Two-pattern and the Related Queries

The *two-pattern query problem* (abbreviated as the 2P problem) was introduced in 2001 by Ferragina et al. [22] and since then it has attracted lots of attention. In the 2P problem, each query is composed of two patterns and a document matches the query if both patterns occur in the document. One can also define the *Forbidden Pattern* (FP) problem [23] where a document matches the query if it contains one pattern but not the other. For symmetry, we also introduce and consider the *Two Forbidden Patterns* (2FP) problem where none of the patterns are allowed to match the document.

Previous Results. Ferragina et al. [22] presented a number of solutions for the 2P problem with space and query times that depend on the “average size” of each document but the worst case query time and space is $\mathcal{O}(P_1 + P_2 + n^\alpha + t)$ and $\mathcal{O}(n^{2-\alpha} \log^{\mathcal{O}(1)} n)$, for any $0 < \alpha < 1$, respectively. Here P_1 and P_2 are the sizes of the query patterns and t is the output size (see also [33]). Cohen and Porat [18] offered a solution that uses $\mathcal{O}(n \log n)$ space with $\mathcal{O}(P_1 + P_2 + \sqrt{nt} \log^{2.5} n)$ query time. The space was improved to $\mathcal{O}(n)$ and the query time to $\mathcal{O}(P_1 + P_2 + \sqrt{nt} \log^{1.5} n)$ by Hon et al. [24]. The query time was reduced by a $\mathcal{O}(\sqrt{\log n})$ in [29] factor and finally the query time $\mathcal{O}(P_1 + P_2 + \sqrt{nt})$ was achieved in [10].

The FP problem was introduced by Fischer et al. [23] and they presented a data structure that stores $\mathcal{O}(n^{3/2})$ bits and answers queries in $\mathcal{O}(P_1 + P_2 + \sqrt{n} + t)$ time. Another solution was given by Hon et al. [25] that uses $\mathcal{O}(n)$ space but has $\mathcal{O}(P_1 + P_2 + \sqrt{nt} \log^{2.5} n)$ query time. For the searching variant (unweighted) their solution can answer queries in

$\mathcal{O}(P_1 + P_2 + \sqrt{n} \log \log n)$ time. As Larsen et al. [29] remark, the $\log \log n$ factor can be removed by using range emptiness data structures and that the same data structure can be used to count the number of matching documents for the two pattern problem.

The difficulty of obtaining fast data structures using (near) linear space has led many to believe that very efficient solutions are impossible to obtain. Larsen et al. [29] specifically focus on proving such impossibility claims and they show that the 2P and FP problems are at least as hard as Boolean Matrix Multiplication, meaning, with current techniques, $P(n) + nQ(n) = \Omega(n^{\omega/2})$ where $P(n)$ and $Q(n)$ are the preprocessing and the query times of the data structure, respectively and ω is the exponent of the best matrix multiplication algorithm (currently $\omega = 2.3728639$). If one assumes that there is no “combinatorial” matrix multiplication algorithm with better running time than $\mathcal{O}(n^{3-o(1)})$, then the lower bound becomes $P(n) + nQ(n) = \Omega(n^{1.5-o(1)})$. Other conditional lower bounds for the 2P and FP problems but from the integer 3SUM conjecture were obtained by Kopelowitz, et al. [27, 28].

The above results are conditional. Furthermore, they tell us nothing about the complexity of the space usage, $S(n)$, versus the query time which is what we are truly interested in for data structure problems. Furthermore, even under the relatively generous assumption¹ that $P(n) = \mathcal{O}(S(n)n^{o(1)})$, the space and query lower bounds obtained from the above results have polynomial gaps compared with the current best data structures.

We need to remark that the only unconditional space lower bound is a pointer machine lower bound that shows with query time of $\mathcal{O}(\text{poly}(\log n) + k)$ the space must be $\Omega(n(\log n / \log \log n)^3)$ [23]. However this bound is very far away from the upper bounds (and also much lower than our lower bounds).

Our Results. We show that all the known data structures for 2P and FP problems are optimal within $n^{o(1)}$ factors, at least in the pointer machine model of computation: Consider a pointer machine data structure that uses $S(n)$ space and can report all the t documents that match a given 2P query (or FP query, or 2FP query) in $Q(n) + \mathcal{O}(P_1 + P_2 + t)$ time. We prove that we must have $S(n)Q(n) = \Omega(n^{2-o(1)})$. As a corollary of our lower bound construction, we also obtain that any data structure that can answer 2P query (or FP query, or 2FP query) in $\mathcal{O}((nt)^{1/2-\varepsilon} + t)$ time, for any fixed constant $\varepsilon > 0$ must use super-linear space. As a side result, we show that surprisingly, the counting variant of the problem is in fact easier: in the semi-group model of computation (see [15] or the full version of this paper [5] for a description of the semi-group model), we prove that we must have

$$S(n)Q^2(n) = \Omega(n^2 / \log^4 n).$$

1.1.2 Set Intersection Queries

The interest in set intersection problems has grown considerably in recent years and variants of the set intersection problem have appeared in many different contexts. Here, we work with the following variants of the problem. The input is m sets, S_1, \dots, S_m of total size n , from a universe \mathcal{U} and a query is a pair of indices i and j . The *decision variant* asks whether $S_i \cap S_j = \emptyset$. The *reporting variant* asks for all the elements in $S_i \cap S_j$. In the *counting variant*, the result should be $|S_i \cap S_j|$. In the *searching variant*, the input also includes a weight function $w : \mathcal{U} \rightarrow G$ where G is a semi-group. The query asks for $\sum_{x \in S_i \cap S_j} w(x)$.

¹ There are problems, such as jumbled indexing [13], where the preprocessing time is a polynomial factor larger than the space complexity.

Previous Results. The set intersection queries have appeared in many different formulations and variants (e.g., see [19, 37, 36, 1, 27, 18]). The most prominent conjecture is that answering the decision variant with constant query time requires $\Omega(n^{2-o(1)})$ space (see [37] for more details). For the reporting variant, Cohen and Porat [18] presented a data structure that uses linear space and answers queries in $\mathcal{O}(\sqrt{nt})$ time, where t is the output size. They also presented a linear-space data structure for the searching variant that answers queries in $\mathcal{O}(\sqrt{n})$ time. In [26] the authors study set intersection queries because of connections to dynamic graph connectivity problems. They offer very similar bounds to those offered by Cohen and Porat (with a $\sqrt{\log n}$ factor worse space and query times) but they allow updates in $\mathcal{O}(\sqrt{n \log n})$ time. It is commonly believed that all set intersection queries are hard. Explicitly stated conjectures on set intersection problems are used to obtain conditional lower bounds for problems such as distance oracles [37, 36, 19] while other well-known conjectures, such as the 3SUM conjecture, can be used to show conditional lower bounds for variants of set intersection problems [27, 1]. For other variants see [9, 34, 21].

Dietz et al. [21] considered set intersection queries in the semi-group model (a.k.a. *the arithmetic model*) and they presented near optimal dynamic and offline lower bounds. They proved that given a sequence of n updates and q queries one must spend $\Omega(q + n\sqrt{q})$ time (ignoring polylog factors); in the offline version a sequence of n insertions and q queries are used but in the dynamic version, the lower bound applies to a dynamic data structure that allows insertion and deletion of points, as well as set intersection queries.

Our Results. Our lower bounds for the 2P problem easily extend to the SI problem. Perhaps the most interesting revelation here is that the searching variant is much easier than the reporting variant ($S(n)Q(n) = \Omega(n^{2-o(1)})$ for reporting versus $S(n)Q(n)^2 = \Omega(n^{2-o(1)})$ for searching)². Based on this, we make another conjecture that even in the RAM model, reporting the elements in $S_i \cap S_j$ for two given query indices i and j , in $\mathcal{O}(n^{1-\varepsilon} + |S_i \cap S_j|)$ time, for any fixed constant $\varepsilon > 0$, requires $\omega(n)$ space. Such a separation between counting and reporting is a rare phenomenon with often counting being the difficult variant.

Observe that conditional lower bounds based on the Boolean Matrix Multiplication or the integer 3SUM conjectures have limitations in distinguishing between the counting and the reporting variants: For example, consider the framework of Larsen et al. [29] and for the best outcome, assume $P(n) = S(n)n^{o(1)}$ and also that boolean matrix multiplication requires $\Omega(n^{3-o(1)})$ time; then their framework yields that $S(n) + nQ(n) = \Omega(n^{3/2-o(1)})$. When $Q(n) = \Theta(n^{2/3})$ this does not rule out the possibility of having $S(n) = \mathcal{O}(n)$ (in fact the counting variant *can* be solved with linear space). However, our lower bound shows that even with $Q(n) = \Theta(n^{2/3})$ the reporting variant requires $\Omega(n^{4/3-o(1)})$ space.

1.1.3 Wild Card Indexing

We study the document indexing problem of matching with “don’t cares”, also known as *wild card matching* or *indexing* (WCI). The setup is the same as for the 2P problem, except a query is a single pattern but it also contains wild cards denoted by a special character “*”. A “*” matches any *one* character. The task is to report all documents matched by the query. This is a well-studied problem from the upper bound perspective and there are generally two variations: either the maximum number of wild cards is bounded or it supports any number

² While we believe our lower bounds are certainly interesting (particularly since they separate counting and reporting variants), they do not make any progress towards resolving the status of the decision version which is considered a major open problem.

of wild cards. We consider the bounded version where patterns contain up to κ wild cards and κ is known in advance by the data structure.

Previous Results. Cole et al. [20] presented a data structure that uses $\mathcal{O}(n \frac{\log^\kappa n}{\kappa!})$ words of space and answers queries in $\mathcal{O}(P + 2^\kappa \log \log n + t)$, where t is the number of occurrences and P is the length of the query pattern. The space has been improved to $\mathcal{O}(n \log^{\kappa+\varepsilon} n)$ bits while keeping the same query time [30]. Another improvement came as a trade-off that increased query time to $\mathcal{O}(P + \beta^j \log \log n + t)$ and reduced space usage $\mathcal{O}(n \log n \log_{\beta}^{\kappa-1} n)$ for any $2 \leq \beta \leq \sigma$ where σ is the alphabet size and $j \leq \kappa$ is the number of wild cards in the query [8]. In the same paper an alternate solution with $\mathcal{O}(P + t)$ query time and $\mathcal{O}(n \sigma^{\kappa^2} \log^\kappa n \log n)$ space usage was also presented. Other results have focused on reducing space while increasing the query time but their query time now depends on the alphabet size, e.g., in [31] the authors provide a data structure with $\mathcal{O}(n \log^\varepsilon n \log \sigma)$ space but with the query time of $\mathcal{O}(m + \sigma^\kappa \sqrt{\log \log \log n} + t)$.

From these bounds we note three things, first that all solutions have some exponential dependency on κ and second the alternate solution by Bille et al. [8] has an odd σ^{κ^2} factor, which is exponential on a *quadratic* function of κ as opposed to being exponential on a linear function of κ (such as 2^κ , σ^κ , or $\log^\kappa n$). Third, when the query time is forced to be independent of σ , there is a discrepancy between query and space when varying κ : Increasing κ (when it is small) by one increases the space by approximately a $\log n$ factor, while to increase the query time by a $\log n$ factor, κ needs to be increased by $\log \log n$. Based on the third point, it is quite likely that unlike SI or 2P problems, WCI does not have a simple trade-off curve.

Other results that are not directly comparable to ours include the following: One is an $\mathcal{O}(n)$ space index with $\mathcal{O}(P + \alpha)$ query time [38] where α is the number of occurrences of all the subpatterns separated by wild card characters. Note that α could be much larger than t and in fact, this can result in a worst case linear query time, even with small values of t . Nonetheless, it could perform reasonably well in practice. Two, there are lower bounds for the *partial match* problem, which is a related problem (see [35] for more details).

Our Results. For WCI with κ wild cards, we prove two results, both in the pointer machine model. In summary, we show that the exponential dependency of space complexity or query time on κ generally cannot be avoided.

As our first result and for a binary alphabet ($\sigma = 2$), we prove that for $3\sqrt{\log n} \leq \kappa = o(\frac{\log n}{\log \log n})$, any data structure that answers queries in $\mathcal{O}(2^{\kappa/2} + P + t)$ time must consume $n^{1+\Theta(1/\log \kappa)}$ space. This result rules out the possibility of lowering the σ^{κ^2} factor in the alternate solution offered by Bille et al. [8], over all values of κ , to $\sigma^{\kappa^{2-\varepsilon}}$ for any constant $\varepsilon > 0$: by setting $\kappa = 3\sqrt{\log n}$ (and $\sigma = 2$), such a bound will be much smaller than our space lower bound (essentially involves comparing $2^{\mathcal{O}(\log^{1-\varepsilon/2} n)}$ factor to a $2^{\Omega(\log n / \log \log n)}$ factor). While this does not rule out improving the space bound for small values of κ , it shows that the exponential dependency on κ^2 is almost tight at least in a particular point in the range of parameters.

As our second result, we prove that answering WCI queries in $Q(n, \kappa) + \mathcal{O}(P + t)$ time requires $\Omega(\frac{n}{\kappa} \Theta(\frac{\log_{Q(n, \kappa)} n}{\kappa})^{\kappa-1})$ space, as long as $\kappa < \log_{Q(n, \kappa)} n$. Note that this query time is assumed to be independent of σ . This result also has a number of consequences. One, it shows that any data structure with query time of $\mathcal{O}(\log^{\mathcal{O}(1)} n + P + t)$ requires $\Omega(\frac{n}{\kappa} (\frac{\log n}{\kappa \log \log n})^{\kappa-1})$ space. Note that this is rather close to the space complexity of the

data structure of Cole et al. [20] ($2^{\mathcal{O}(\kappa)}(\log \log n)^\kappa$ factor away). In other words, the data structure of Cole et al. cannot be significantly improved both in space complexity and query time; e.g., it is impossible to answer WCI queries in $\mathcal{O}(\log^{\mathcal{O}(1)} n + P + t)$ time using $\mathcal{O}\left(n \left(\frac{\log n}{\kappa}\right)^{\kappa(1-\varepsilon)}\right)$ space, for any constant $\varepsilon > 0$.

Two, $\kappa = 2$ is the smallest value where linear space becomes impossible with polylogarithmic query time. This is very nice since $\kappa = 1$ can be solved with almost linear space [30]. Furthermore this shows that the increase by a $\log n$ factor in the space complexity for every increase of κ is necessary (for small values of κ).

Three, we can combine our second result with our first result when $\kappa = 3\sqrt{\log n}$. As discussed before, our first result rules out fast queries (e.g., when $Q(n) \leq 2^{\kappa/2}$), unless the data structure uses large amounts of space, so consider the case when $Q(n) = \mathcal{O}(2^\kappa)$. In this case, we can rule out lowering the space usage of the data structure of Cole et al. to $\Omega\left(n \left(\frac{\log n}{\kappa}\right)^{\kappa^{1-\varepsilon}}\right)$ for any constant $\varepsilon > 0$: apply our second lower bound with fewer wild cards, specifically, with $\kappa' = \kappa^{1-\delta}$ wild cards, for a small enough constant $\delta > 0$ that depends on ε . Observe that $\kappa' < \log_{Q(n)} n$, so the second result lower bounds the space by $\Omega\left(\frac{n}{\kappa'} \Theta\left(\log^{\frac{\delta}{2}\kappa^{1-\delta}} n\right)\right)$, which for a sufficiently small δ is greater than $\Omega\left(n \left(\frac{\log n}{\kappa}\right)^{\kappa^{1-\varepsilon}}\right)$.

As mentioned in the beginning, our results show that the exponential dependency of space or query time on κ cannot be improved in general. Furthermore, at a particular point in the range of parameters (when $\kappa = 3\sqrt{\log n}$), all the known exponential dependencies on κ are almost tight and cannot be lowered to an exponential dependency on $\kappa^{1-\varepsilon}$ (or on $\kappa^{2-\varepsilon}$ for the alternate solution) for any constant $\varepsilon > 0$. Nonetheless, there are still gaps between our lower bounds and the known data structures. We believe it is quite likely both our lower bounds and the existing data structures can be improved to narrow the gap.

1.1.4 Gapped Pattern Indexing

A κ -gapped pattern is a pattern $p_1\{\alpha_1, \beta_1\}p_2\{\alpha_2, \beta_2\}, \dots, p_\kappa\{\alpha_\kappa, \beta_\kappa\}p_{\kappa+1}$ where α_i and β_i , $1 \leq i \leq \kappa$ are integers, and each p_i , $1 \leq i \leq \kappa + 1$, is a string over an alphabet of size σ . Such a κ -gapped pattern matches a documents in which one can find one occurrence of every p_i such that there are at least α_i and at most β_i characters between the occurrence of p_i and the occurrence of p_{i+1} , $1 \leq i \leq \kappa$.

Previous Results. The gapped pattern indexing is often considered both in the online and the offline version (e.g., see [7, 40]). However, the result most relevant to us is [6], where they consider the following data structure problem: given a set of 1-gapped patterns of total size n , where all the patterns are in the form of $p_1\{\alpha, \beta\}p_2$, store them in a data structure such that given a document of length D at the query time, one can find all the gapped patterns that match the query document (in general we call this the κ -gapped pattern indexing (κ -GPI) problem). In [6], the authors give a number of upper bounds and conditional lower bounds for the problem. Among a number of results, they can build a data structure of linear size that can answer queries in $\tilde{\mathcal{O}}(D(\beta - \alpha) + t)$ where $\tilde{\mathcal{O}}$ notation hides polylogarithmic factors and t is the output size. For the lower bound and among a number of results, they can show that with linear space $\Omega(D(\beta - \alpha)^{1-o(1)} + t)$ query time is needed.

Our Results. We consider the general κ -GPI problem where $\beta_i - \alpha_i = \gamma$ for all $1 \leq i \leq \kappa$, and a prove lower bound that is surprisingly very high: any pointer machine data structure

that can answer queries in $o(D\gamma^\kappa/(2\log D)^{\kappa+1})$ time must use *super polynomial* space of $n^{\Omega(\log^{1/(\kappa+1)} n)}$. By construction, this result also holds if the input is a set of $\kappa + 1$ patterns where they all need to match the query document (regardless of their order and size of the gaps). In this case, answering queries in $o(D^{\kappa+1}/(2\log D)^{\kappa+1})$ requires the same super-polynomial space. Note that in this case $\kappa = 1$ is the “dual” of the 2P problem: store a set of two-patterns in data structure such that given a query document, we can output the subset of two-patterns that match the document.

1.2 Technical Preliminaries

The Pointer Machine Model [39]. This models data structures that solely use pointers to access memory locations (e.g., any tree-based data structure)³. We focus on a variant that is the popular choice when proving lower bounds [14]. Consider an abstract “reporting” problem where the input includes a universe set \mathcal{U} where each query q reports a subset, $q_{\mathcal{U}}$, of \mathcal{U} . The data structure is modelled as a directed graph G with outdegree two (and a root node $r(G)$) where each vertex represents one memory cell and each memory cell can store one element of \mathcal{U} ; edges between vertices represent pointers between the memory cells. All other information can be stored and accessed for free by the data structure. The only requirement is that given the query q , the data structure must start at $r(G)$ and explore a connected subgraph of G and find its way to vertices of G that store the elements of $q_{\mathcal{U}}$. The size of the subgraph explored is a lower bound on the query time and the size of G is a lower bound on the space.

An important remark. The pointer-machine can be used to prove lower bounds for data structures with query time $Q(n) + \mathcal{O}(t)$ where t is the output size and $Q(n)$ is “search overhead”. Since we can simulate any RAM algorithm on a pointer-machine with $\log n$ factor slow down, we cannot hope to get high unconditional lower bounds if we assume the query time is $Q(n) + \mathcal{O}(t \log n)$, since that would automatically imply RAM lower bounds for data structures with $Q(n)/\log n + \mathcal{O}(t)$ query time, something that is hopelessly impossible with current techniques. However, when restricted to query time of $Q(n) + \mathcal{O}(t)$, the pointer-machine model is an attractive choice and it has an impressive track record of proving lower bounds that match the best known data structures up to very small factors, even when compared to RAM data structures; we mention two prominent examples here. For the fundamental *simplex range reporting* problem, all known solutions are pointer-machine data structures [32, 11, 17] and the known pointer machine lower bounds match these up to an $n^{o(1)}$ factor [2, 16]. One can argue that it is difficult to use the power of RAM for simplex range reporting problem. However, for the other fundamental *orthogonal range reporting*, where it is easy to do various RAM tricks, the best RAM data structures save at most a $\log n$ factor compared to the best known pointer machine solutions (e.g., see [3, 4, 12]). Also, where cell-probe lower bounds cannot break the $\log n$ query barrier, very high lower bounds are known for the orthogonal range reporting problem in the pointer machine model [3, 4, 14].

Known Frameworks. The fundamental limitation in the pointer machine model is that starting from a memory cell v , one can visit at most 2^ℓ other memory cells using ℓ pointer

³ Many of the known solutions for various indexing problems use tree structures, such as suffix trees or wavelet trees. While sometimes trees can be encoded using bit-vectors with rank/select structures on top, these tricks can only save polylogarithmic factors in space and query times.

navigations. There are two known methods that exploit this limitation and build two different frameworks for proving lower bounds.

The first lower bound framework was given by Bernard Chazelle [14, 16]. However, we will need a slightly improved version of his framework that is presented in the following lemma; essentially, we need a slightly tighter analysis on a parameter that was originally intended as a large constant. Due to lack of space, the proof is only available in the full version [5].

► **Theorem 1.** *Let \mathcal{U} be a set of n input elements and \mathcal{Q} a set of queries where each query outputs a subset of \mathcal{U} . Assume there exists a data structure that uses $S(n)$ space and answers each query in $Q(n) + \alpha k$ time, where k is the output size. Assume (i) the output size of any query $q \in \mathcal{Q}$, denoted by $|\mathcal{U} \cap q|$, is at least t , for a parameter $t \geq Q(n)$ and (ii) for integers ℓ and β , and indices, i_1, \dots, i_ℓ , $|\mathcal{U} \cap q_{i_1} \cap \dots \cap q_{i_\ell}| < \beta$. Then, $S(n) = \Omega\left(\frac{|\mathcal{Q}|t}{\ell \cdot 2^{O(\alpha\beta)}}\right)$.*

The second framework is due to Afshani [2] and it is designed for “geometric stabbing problems”: given an input set of n geometric regions, the goal is store them in a data structure such that given a query point q , one can output the subset of regions that contain q . The framework is summarized below.

► **Theorem 2 ([2]).** *Assume one can construct a set of n geometric regions inside the d -dimensional unit cube such that (i) every point of the unit cube is contained in at least t regions⁴, and (ii) the volume of the intersection of every β regions is at most v , for some parameters β , t , and v . Then, for any pointer-machine data structure that uses $S(n)$ space and can answer geometric stabbing queries on the above input in time $g(n) + O(k)$, where k is the output size and $g(\cdot)$ is some increasing function, if $g(n) \leq t$ then $S(n) = \Omega(tv^{-1}2^{-O(\beta)})$.*

These two frameworks are not easily comparable. In fact, for many constructions, often only one of them gives a non-trivial lower bound. Furthermore, as remarked by Afshani [2], Theorem 2 does not need to be operated in the d -dimensional unit cube and in fact any measure could be substituted instead of the d -dimensional Lebesgue measure.

Our Techniques. Our first technical contribution is to use Theorem 2 in a non-geometric setting by representing queries as abstract points under a discrete measure and each input object as a range that contains all the matching query points. Our lower bound for the κ -GPI problem and one of our WCI lower bounds are proved in this way. The second technical contribution is actually building and analyzing proper input and query sets to be used in the lower bound frameworks. In general, this is not easy and in fact for some problems it is highly challenging⁵. Also see Section 5 (Conclusions) for a list of open problems.

In the rest of this article, we present the technical details behind our κ -GPI lower bound and most of the details of our first WCI lower bound. Due to lack of space, the rest of the technical details can be found in the full version [5].

We begin with the κ -GPI problem since it turns out for this particular problem we can get away with a simple deterministic construction. For WCI, we need a more complicated randomized construction to get the best result and thus it is presented next.

⁴ In [2] this is stated as “exactly t ranges” but the proof works with only a lower bound on t .

⁵ A good example is the classical halfspace range reporting problem where constructing proper input and query sets has been a longstanding open problem; the current best lower bound uses a highly inefficient reduction from the simple range reporting problem [2].

2 Gapped Pattern Lower Bound

In this section we deal with the data structure version of the gapped pattern problem. The input is a collection of κ -gapped patterns (typically called a dictionary), with total length n (in characters). The goal is to store the input in a data structure such that given a document of size D , one can report all the input gapped patterns that match the query. We focus on special κ -gapped patterns that we call *standard*: a standard κ -gapped pattern in the form of $p_1 \{0, \gamma\} p_2 \{0, \gamma\} \dots \{0, \gamma\} p_{\kappa+1}$ where each p_i is a string (which we call a *subpattern*) and γ is an integer.

► **Theorem 3.** For $\kappa = o\left(\frac{\log \log n}{\log \log \log n}\right)$ and in the pointer machine model, answering κ -GPI queries in $o\left(\frac{D\gamma^\kappa}{(2+2\log D)^{\kappa+1}}\right) + O(t)$ time requires $n^{\Omega(\log^{1/(2\kappa)} n)}$ space.

To prove this theorem, we build a particular input set of standard κ -gapped patterns. We pick the alphabet $\Sigma = \{0, 1, \#\}$, and the gapped patterns only use $\{0, 1\}$. Each subpattern in the input is a binary string of length p . The subpatterns in any gapped pattern are in lexicographic order, and a subpattern occurs at most once in a pattern (i.e., no two subpatterns in a pattern are identical). The input set, S , contains all possible gapped patterns obeying these conditions. Thus, $|S| = \binom{2^p}{\kappa+1}$. Each query is a text composed of concatenation of $D/(p+1)$ substrings (for simplicity, we assume $D/(p+1)$ is an integer) and each substring is in the form $\#\{0, 1\}^p$. We restrict the query substrings to be in lexicographic order and without repetitions (no two substrings in a query are identical). The set of all query texts satisfying these constraints is denoted by Q . Thus, $|Q| = \binom{2^p}{D/(p+1)}$.

► **Lemma 4.** For $D \geq 2\kappa\gamma$, any text $T \in Q$ matches $\Theta\left(\frac{D\gamma^\kappa}{(p+1)^{\kappa+1}}\right)$ gapped patterns in S .

Proof. Consider a text $T \in Q$. To count the number of gapped patterns that can match it, we count the different ways of selecting $\kappa+1$ positions that correspond to the starting positions of a matching subpattern. Each position starts immediately after $\#$, with at most γ characters between consecutive positions. Since $D \geq 2\kappa\gamma$, we have $\Theta(D/p)$ choices for picking the first position, i.e., the starting position of a gapped pattern matching T . After fixing the first match, there are at most $\gamma/(p+1)$ choices for the position of the next match between a subpattern and substring. However, if the first match happens in the first half of text T , there are always $\gamma/(p+1)$ choices for the position of each subpattern match (since $D \geq 2\kappa\gamma$). Thus, we have $\Theta(D\gamma^\kappa/(p+1)^{\kappa+1})$ choices. As input subpatterns are in lexicographically increasing order, different choices result in distinct gapped patterns that match the query. ◀

To apply Theorem 2, we consider each query text T in Q as a “discrete” point with measure $1/|Q|$. Thus, the total measure of Q (i.e., the query space) is one and Q functions as the “unit cube” within the framework of Theorem 2. We consider an input gapped pattern P in S as a range that contains all the points of Q that match P . Thus, to apply Theorem 2, we need to find a lower bound on the output size of every query (condition (i)) and an upper bound on v , the measure of the intersection of β inputs (condition (ii)). By the above lemma, meeting the first condition is quite easy: we pick $t = \Theta\left(\frac{D\gamma^\kappa}{(p+1)^{\kappa+1}}\right)$ (with the right hidden constant). Later we shall see that $p = 1 + 2\log D$ so this can be written as $t = \Theta\left(\frac{D\gamma^\kappa}{(2+2\log D)^{\kappa+1}}\right)$. Thus, we only need to upper bound v which we do below.

► **Lemma 5.** Consider β patterns $P_1, \dots, P_\beta \in S$. At most $\binom{2^p - \beta^{1/(\kappa+1)}}{D/(p+1) - \beta^{1/(\kappa+1)}}$ texts in Q can match all patterns P_1, \dots, P_β .

Proof. Collectively, P_1, \dots, P_β must contain at least $r = \beta^{1/(\kappa+1)}$ distinct subpatterns: otherwise, we can form at most $\binom{r}{\kappa+1} < \beta$ different gapped patterns, a contradiction. This in turn implies that any text T matching P_1, \dots, P_β must contain all these at least r distinct subpatterns. Clearly, the number of such texts is at most $\binom{2^p - \beta^{1/(\kappa+1)}}{D/(p+1) - \beta^{1/(\kappa+1)}}$. ◀

As the measure of each query in Q is $1/|Q|$, by the above theorem, we have $v \leq \binom{2^p - \beta^{1/(\kappa+1)}}{D/(p+1) - \beta^{1/(\kappa+1)}}/|Q|$. We now can apply Theorem 2. Each pattern in the input has $\Theta(p\kappa)$ characters and thus the total input size, n , is $\Theta(p\kappa|S|) = \Theta(p\kappa \binom{2^p}{\kappa+1})$. By the framework, and Lemmata 5 and 4, we know that the space usage of the data structure is at least

$$\Omega \left(\frac{D\gamma^\kappa}{(p+1)^{\kappa+1}} \cdot \frac{\binom{2^p}{D/(p+1)}}{\binom{2^p - \beta^{1/(\kappa+1)}}{D/(p+1) - \beta^{1/(\kappa+1)}}} \cdot 2^{-\mathcal{O}(\beta)} \right) = \Omega \left(1 \cdot \left(\frac{(p+1)2^{(p-1)}}{D} \right)^{\beta^{1/(\kappa+1)}} \cdot 2^{-\mathcal{O}(\beta)} \right)$$

where to obtain the rightmost equation we expand the binomials, simplify, and constrain $\beta^{1/(\kappa+1)} < 2^p/2$ to lower bound $2^p - \beta^{1/(\kappa+1)}$ with 2^{p-1} . Now we work out the parameters. We know that $n = \Theta(p\kappa \binom{2^p}{\kappa+1}) = 2^{\Theta(p(\kappa+1))}$; this is satisfied by setting $p = c_p(\log n)/(\kappa+1)$ for some constant c_p . Observe that there is an implicit constraint on D and p : there should be sufficient bits in the subpatterns to fill out a query with distinct subpatterns, i.e. $2^p > D/(p+1)$; we pick $D = n^{c_D 1/(\kappa+1)}$ for some other constant c_D such that $D = 2^{p/2-1}$ and thus $(p+1)2^{p-1}/D = 2^{p/2}$. Using these values, the space lower bound is simplified to

$$\Omega \left(2^{\beta^{1/(\kappa+1)} \frac{c_p \log n}{2} \frac{1}{\kappa+1}} \cdot 2^{-c_\beta \beta} \right)$$

where c_β is another constant. We now optimize the lower bound by picking β such that $c_\beta \beta = \frac{1}{2} \beta^{1/(\kappa+1)} \frac{c_p \log n}{2} \frac{1}{\kappa+1}$ which solves to $\beta = \Theta\left(\left(\frac{\log n}{\kappa+1}\right)^{1+1/\kappa}\right)$. Thus, for constant c , the space complexity of the data structure is

$$\Omega \left(2^{c \left(\frac{\log n}{\kappa+1}\right)^{1+\frac{1}{\kappa}}} \right) = \Omega \left(2^{\log n \cdot c \log^{\frac{1}{\kappa}} n \left(\frac{1}{\kappa+1}\right)^{1+\frac{1}{\kappa}}} \right) = n^{\Omega \left(\log^{\frac{1}{\kappa}} n \left(\frac{1}{\kappa+1}\right)^{1+\frac{1}{\kappa}} \right)} = n^{\Omega \left(\log^{\frac{1}{2\kappa}} n \right)}$$

where the last part follows from $\kappa = o(\log \log n / \log \log \log n)$.

3 Wild Card Indexing

In this section we consider the wild card indexing (WCI) problem and prove both space and query lower bounds in the pointer machine model of computation. Note that our query lower bound applies to an alphabet size of two (i.e., binary strings).

3.1 The Query Lower Bound

Assume for any input set of documents of total size n , we can build a data structure such that given a WCI query of length m containing κ wild cards, we can find all the documents that match the query in $Q(n, \kappa) + O(m+t)$ time, where t is the output size. Furthermore, assume κ is a fixed parameter known by the data structure and that the data structure consumes $S(n, \kappa)$ space. Our main result here is the following.

► **Theorem 6.** *If $3\sqrt{\log n} \leq \kappa = o(\log n)$ and $Q(n, \kappa) = O(2^{\frac{\kappa}{2}})$, then $S(n, \kappa) \geq n^{1+\Theta(\frac{1}{\log \kappa})}$.*

To prove the lower bound, we build a *particular* set of documents and patterns and prove that if the data structure can answer the queries fast, then it must consume lots of space, for this particular input, meaning, we get lower bounds for the function $S(\cdot)$. We now present the details. We assume $Q(n, \kappa) \leq 2^{\kappa/2}$, as otherwise the theorem is trivial.

Documents and patterns. We build the set of documents in two stages. Consider the set of all bit strings of length m with exactly $\ell = \kappa/2$ “1”s. In the first stage, we sample each such string independently with probability r^{-1} where $r = 2^{\kappa/3}$. Let \mathcal{D} be the set of sampled strings. In the second stage, for every set of $\ell + \ell'$ indices, $1 \leq i_1 < i_2 < \dots < i_{\ell+\ell'} \leq m$, where $\ell' = (\log_\ell r)/2 = \Theta(\kappa/\log \kappa)$, we perform the following operation, given another parameter β : if there are more than β strings in \mathcal{D} that have “1”s only among positions $i_1, \dots, i_{\ell+\ell'}$, then we remove all such strings from \mathcal{D} . Consequently, among the remaining strings in \mathcal{D} , “1”s in every subset of β strings will be spread over at least $\ell + \ell' + 1$ positions. The set of remaining strings \mathcal{D} will form our input set of documents. Now we consider the set \mathcal{P} of all the patterns of length m that have exactly κ wild cards and $m - \kappa$ “0”s. We remove from \mathcal{P} any pattern that matches fewer than $\binom{\kappa}{\ell}/(2r)$ documents from \mathcal{D} . The remaining patterns in \mathcal{P} will form our query set of patterns. In the full version [5], we prove the following.

► **Lemma 7.** *With positive probability, we get a set \mathcal{D} of $\Theta(\binom{m}{\ell}/r)$ documents and a set \mathcal{P} of $\Theta(\binom{m}{\kappa})$ patterns such that (i) each pattern matches $\Theta(\binom{\kappa}{\ell}/r)$ documents, and (ii) there are no $\beta = \Theta(\log_\kappa m)$ documents whose “1”s are contained in a set of $\ell + \ell'$ indices.*

To prove the lower bound, we use Theorem 2. We use a discrete measure here: each pattern is modelled as a “discrete” point with measure $\frac{1}{|\mathcal{P}|}$, meaning, the space of all patterns has measure one. Each document forms a range: a document d_i contains all the patterns (discrete points) that match d_i . Thus, the measure of every document $d_i \in \mathcal{D}$ is $t_i/|\mathcal{P}|$, where t_i is the number of patterns that match d_i . We consider the measure of the intersection of β documents (regions) d_1, \dots, d_β . By Lemma 7, there are $\ell + \ell'$ indices where one of these documents has a “1”; any pattern that matches all of these documents must have a wild card in all of those positions. This means, there are at most $\binom{m-\ell-\ell'}{\kappa-\ell-\ell'}$ patterns that could match documents d_1, \dots, d_β . This means, when we consider documents as ranges, the intersection of every β documents has measure at most $\binom{m-\ell-\ell'}{\kappa-\ell-\ell'}/|\mathcal{P}|$ which is an upper bound for parameter v in Theorem 2. For the two other parameters t and $g(n)$ in the theorem we have, $t = \Theta(\binom{\kappa}{\ell}/r)$ (by Lemma 7) and $g(n) = Q(n, \kappa) + O(m)$. To obtain a space lower bound from Theorem 2, we must check if $g(n) \leq t = \Theta(\binom{\kappa}{\ell}/r)$. Observe that $\binom{\kappa}{\ell} = \binom{\kappa}{\kappa/2} \geq 2^\kappa/\kappa$ since the binomials $\binom{\kappa}{i}$ sum to 2^κ for $0 \leq i \leq \kappa$ and $\binom{\kappa}{\kappa/2}$ is the largest one. As $r = 2^{-\kappa/3}$, we have $t = \Omega(2^\kappa 2^{-\kappa/3}/\kappa) = \omega(2^{\kappa/2}) = \omega(Q(n, \kappa))$. However, $g(n)$ also involves an additive $O(m)$ term. Thus, we must also have $t = \omega(m)$ which will hold for our choice of parameters but we will verify it later.

By guaranteeing that $g(n) \leq t$, Theorem 2 gives a space lower bound of $\Omega(tv^{-1}2^{-O(\beta)})$. However, we would like to create an input of size $\Theta(n)$ which means the number of sampled documents must be $\Theta(n/m)$ and thus we must have $\binom{m}{\ell}/r = \Theta(n/m)$. As $m = \omega(\kappa)$, it follows that $\binom{m}{\ell} = (\Theta(1)m/\ell)^\ell$. Thus, $(\Theta(1)m/\ell)^\ell = \Theta(rn/m)$. Thus, we have that

$$v^{-1} \geq \frac{|P|}{\binom{m-\ell-\ell'}{\kappa-\ell-\ell'}} \geq \frac{\Theta\left(\binom{m}{\kappa}\right)}{\binom{m-\ell-\ell'}{\kappa-\ell-\ell'}} = \Theta(1) \cdot \frac{\frac{m!}{\kappa!(m-\kappa)!}}{\frac{(m-\ell-\ell')!}{(\kappa-\ell-\ell')!(m-\kappa)!}} \geq \Theta(1) \cdot \frac{m^{\ell+\ell'}}{(2\kappa)^{\ell+\ell'}} = \frac{\Theta(1) \cdot n}{m\Theta(1)^\kappa} \cdot \left(\frac{m}{2\kappa}\right)^\ell$$

where the last step follows from $(\Theta(1)m/\ell)^\ell = \Theta(rn/m)$, $r = 2^{\kappa/3}$ and $\ell = \kappa/2$.

Now we bound m in terms of n . From $(m/(2\kappa))^\ell = \frac{n}{m\Theta(1)^\kappa}$ we obtain that $m = \kappa^{\ell/(\ell+1)} \cdot n^{1/(\ell+1)}/\Theta(1)^\kappa = n^{2/(\kappa+2)}/\Theta(1)^\kappa$. Remember that $\ell' = \Theta(\kappa/\log \kappa)$. Based on this, we get that $v^{-1} = n \cdot n^{\Theta(1/\log \kappa)}/\Theta(1)^\kappa$ and since $\kappa = o(\log n/\log \log n)$ the $\Theta(1)^\kappa$ term is dominated and we have $v^{-1} = n \cdot n^{\Theta(1/\log \kappa)}$. It remains to handle the extra $2^{-O(\beta)}$ factor in the space lower bound. From Lemma 7, we know that $\beta = c \log_\kappa m$. Based on the value of m , this means $\beta = \Theta(\log n/(\kappa \log \kappa))$ which means $2^{-O(\beta)}$ is also absorbed in $n^{\Theta(1/\log \kappa)}$

factor. It remains to verify one last thing: previously, we claimed that we would verify that $t = \omega(m)$. Using the bound $t = \omega(2^{\kappa/2})$ this can be written as $2^{\kappa/2} = \omega(m)$ which translates to $\kappa/2 = (2 \log n)/(\kappa + 2) + \omega(1)$ which clearly holds if $\kappa \geq 3\sqrt{\log n}$.

3.2 The Space Lower Bound

The details of our space lower can be found in the full version [5], where we prove the following.

► **Theorem 8.** *labelthm:wci* Any pointer-machine data structure that answers WCI queries with κ wild cards in time $Q(n) + O(m + t)$ over an input of size n must use $\Omega\left(\frac{n}{\kappa} \Theta\left(\frac{\log_{Q(n)} n}{\kappa}\right)^{\kappa-1}\right)$ space, as long as $\kappa < \log_{Q(n)} n$, where t is the output size, and m is the pattern length.

Refer to the introduction for a discussion of the consequences of these lower bounds.

4 Two Pattern Document Indexing and Related Problems

Due to lack of space we only state the primary results for 2P, FP, 2FP, and SI. The proofs for Theorems 9 and 10 can be found in the full version [5].

► **Theorem 9.** Any data structure on the Pointer Machine for the 2P, FP, 2FP, and SI problems with query time $Q(n)$ and space usage $S(n)$ must obey $S(n)Q(n) = \Omega(n^{2-o(1)})$.

Also, if $Q(n, k) = O((nk)^{1/2-\alpha} + k)$ for a constant $0 < \alpha < 1/2$, then $S(n) = \Omega\left(n^{\frac{1+6\alpha}{1+2\alpha}-o(1)}\right)$.

The above theorem is proved using Theorem 1 which necessitates a randomized construction involving various high probability bounds. Unlike our lower bound for κ -GPI we were unable to find a deterministic construction that uses Theorem 2.

We also prove the following lower bound in the semi-group model which addresses the difficulty of the counting variants of 2P and the related problems.

► **Theorem 10.** Answering 2P, FP, 2FP, and SI queries in the semi-group model requires $S(n)Q^2(n) = \Omega(n^2/\log^4 n)$.

5 Conclusions

In this paper we proved unconditional and high space and query lower bounds for a number of problems in string indexing. Our main message is that the pointer machine model remains an extremely useful tool for proving lower bounds, that are close to the *true complexity* of many problems. We have successfully demonstrated this fact in the area of string and document indexing. Within the landscape of lower bound techniques, the pointer machine model, fortunately or unfortunately, is the only model where we can achieve *unconditional, versatile, and high* lower bounds and we believe more problems from the area of string and document indexing deserve to be considered in this model. To this end, we outline a number of open problems connected to our results.

1. Is it possible to generalize the lower bound for 2P to the case where the two patterns are required to match within distance γ ? This is essentially a "dual" of the 1-GPI problem.
2. Recall that our space lower bound for the WCI problem (Subsection 3.2) assumes that the query time is independent of the alphabet size σ . What if the query is allowed to increase with σ ?

3. Our query lower bound for the WCI (Subsection 3.1) is proved for a binary alphabet. Is it possible to prove lower bounds that take σ into account? Intuitively the problem should become more difficult as σ increases, but we were unable to obtain such bounds.
4. We require certain bounds on κ for the WCI problem. Is it possible to remove or at least loosen them? Or perhaps, can the upper bounds be substantially improved?
5. What is the difficulty of the κ -GPI problem when κ is large?

Acknowledgements. We thank Karl Bringmann for simplifying the construction of our documents for the 2P problem. We thank Moshe Lewenstein for bringing the WCI problem to our attention. We also thank the anonymous referees for pointing us to [21].

References

- 1 Amir Abboud and Virginia Vassilevska Williams. Popular conjectures imply strong lower bounds for dynamic problems. In *Proceedings of Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 434–443, 2014.
- 2 Peyman Afshani. Improved pointer machine and I/O lower bounds for simplex range reporting and related problems. In *Symposium on Computational Geometry (SoCG)*, pages 339–346, 2012.
- 3 Peyman Afshani, Lars Arge, and Kasper Dalgaard Larsen. Orthogonal range reporting: query lower bounds, optimal structures in 3-d, and higher-dimensional improvements. In *Symposium on Computational Geometry (SoCG)*, pages 240–246, 2010.
- 4 Peyman Afshani, Lars Arge, and Kasper Green Larsen. Higher-dimensional orthogonal range reporting and rectangle stabbing in the pointer machine model. In *Symposium on Computational Geometry (SoCG)*, pages 323–332, 2012. doi:10.1145/2261250.2261299.
- 5 Peyman Afshani and Jesper Sindahl Nielsen. Data structure lower bounds for document indexing problems. *CoRR*, abs/1604.06264, 2016. URL: <http://arxiv.org/abs/1604.06264>.
- 6 Amihoud Amir, Tsvi Kopelowitz, Avivit Levy, Seth Pettie, Ely Porat, and B. Riva Shalom. Online Dictionary Matching with One Gap. *CoRR*, abs/1503.07563, 2015. URL: <http://arxiv.org/abs/1503.07563>.
- 7 Amihoud Amir, Avivit Levy, Ely Porat, and B. Riva Shalom. Dictionary Matching with One Gap. In *Annual Symposium on Combinatorial Pattern Matching (CPM)*, pages 11–20, 2014.
- 8 Philip Bille, Inge Li Gørtz, Hjalte Wedel Vildhøj, and Søren Vind. String indexing for patterns with wildcards. In *Scandinavian Symposium and Workshops on Algorithm Theory (SWAT)*, pages 283–294, 2012.
- 9 Philip Bille, Anna Pagh, and Rasmus Pagh. Fast evaluation of union-intersection expressions. In *Algorithms and Computation*, volume 4835 of *Lecture Notes in Computer Science*, pages 739–750. Springer Berlin Heidelberg, 2007.
- 10 Sudip Biswas, Arnab Ganguly, Rahul Shah, and Sharma V. Thankachan. Ranked Document Retrieval with Forbidden Pattern. In *Annual Symposium on Combinatorial Pattern Matching (CPM)*, pages 77–88, 2015.
- 11 Timothy M. Chan. Optimal partition trees. In *Symposium on Computational Geometry (SoCG)*, pages 1–10. ACM, 2010.
- 12 Timothy M. Chan, Kasper Green Larsen, and Mihai Pătraşcu. Orthogonal range searching on the RAM, revisited. In *Symposium on Computational Geometry (SoCG)*, pages 1–10, 2011.

- 13 Timothy M. Chan and Moshe Lewenstein. Clustered integer 3SUM via additive combinatorics. In *Proceedings of ACM Symposium on Theory of Computing (STOC)*, pages 31–40, 2015.
- 14 Bernard Chazelle. Lower Bounds for Orthogonal Range Searching: I. The Reporting Case. *J. ACM*, 37(2):200–212, 1990.
- 15 Bernard Chazelle. Lower Bounds for Orthogonal Range Searching II. The Arithmetic Model. *J. ACM*, 37(3):439–463, 1990.
- 16 Bernard Chazelle and Burton Rosenberg. Simplex Range Reporting on a Pointer Machine. *Comput. Geom.*, 5:237–247, 1995.
- 17 Bernard Chazelle, Micha Sharir, and Emo Welzl. Quasi-optimal upper bounds for simplex range searching and new zone theorems. *Algorithmica*, 8:407–429, December 1992.
- 18 Hagai Cohen and Ely Porat. Fast set intersection and two-patterns matching. *Theor. Comput. Sci.*, 411(40-42):3795–3800, 2010.
- 19 Hagai Cohen and Ely Porat. On the hardness of distance oracle for sparse graph. <http://arxiv.org/abs/1006.1117>, 2010.
- 20 Richard Cole, Lee-Ad Gottlieb, and Moshe Lewenstein. Dictionary matching and indexing with errors and don't cares. In *Proceedings of ACM Symposium on Theory of Computing (STOC)*, pages 91–100, 2004.
- 21 Paul F. Dietz, Kurt Mehlhorn, Rajeev Raman, and Christian Uhrig. Lower bounds for set intersection queries. *Algorithmica*, 14(2):154–168, 1995.
- 22 Paolo Ferragina, Nick Koudas, S. Muthukrishnan, and Divesh Srivastava. Two-dimensional substring indexing. *Journal of Computer and System Sciences*, 66(4):763–774, 2003. Special Issue on PODS 2001.
- 23 Johannes Fischer, Travis Gagie, Tsvi Kopelowitz, Moshe Lewenstein, Veli Mäkinen, Leena Salmela, and Niko Välimäki. Forbidden patterns. In *LATIN 2012: Theoretical Informatics – 10th Latin American Symposium, Arequipa, Peru, April 16-20, 2012. Proceedings*, pages 327–337, 2012.
- 24 Wing-Kai Hon, Rahul Shah, Sharma V. Thankachan, and Jeffrey Scott Vitter. String Retrieval for Multi-pattern Queries. In *String Processing and Information Retrieval – 17th International Symposium, SPIRE 2010, Los Cabos, Mexico, October 11-13, 2010. Proceedings*, pages 55–66, 2010.
- 25 Wing-Kai Hon, Rahul Shah, Sharma V. Thankachan, and Jeffrey Scott Vitter. Document Listing for Queries with Excluded Pattern. In *Combinatorial Pattern Matching – 23rd Annual Symposium, CPM 2012, Helsinki, Finland, July 3-5, 2012. Proceedings*, pages 185–195, 2012.
- 26 Casper Kejlberg-Rasmussen, Tsvi Kopelowitz, Seth Pettie, and Ely Porat. Word-packing algorithms for dynamic connectivity and dynamic sets. <http://arxiv.org/abs/1407.6755>, 2014.
- 27 Tsvi Kopelowitz, Seth Pettie, and Ely Porat. Higher lower bounds from the 3SUM conjecture. <http://arxiv.org/abs/1407.6756>, 2016.
- 28 Tsvi Kopelowitz, Seth Pettie, and Ely Porat. Higher lower bounds from the 3SUM conjecture. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1272–1287, 2016.
- 29 Kasper Green Larsen, J. Ian Munro, Jesper Sindahl Nielsen, and Sharma V. Thankachan. On Hardness of Several String Indexing Problems. In *Combinatorial Pattern Matching – 25th Annual Symposium, CPM 2014, Moscow, Russia, June 16-18, 2014. Proceedings*, pages 242–251, 2014.
- 30 Moshe Lewenstein, J. Ian Munro, Venkatesh Raman, and Sharma V. Thankachan. Less space: Indexing for queries with wildcards. *Theoretical Computer Science*, 557:120–127, 2014.

- 31 Moshe Lewenstein, Yakov Nekrich, and Jeffrey Scott Vitter. Space-efficient string indexing for wildcard pattern matching. *CoRR*, abs/1401.0625, 2014. URL: <http://arxiv.org/abs/1401.0625>.
- 32 Jiří Matoušek. Range searching with efficient hierarchical cuttings. *Discrete & Computational Geometry*, 10(2):157–182, 1993.
- 33 S. Muthukrishnan. Efficient algorithms for document retrieval problems. In *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 657–666, 2002.
- 34 Mihai Pătraşcu. Towards polynomial lower bounds for dynamic problems. In *Proceedings of ACM Symposium on Theory of Computing (STOC)*, pages 603–610, 2010.
- 35 Mihai Pătraşcu. Unifying the landscape of cell-probe lower bounds. *SIAM Journal of Computing*, 40(3):827–847, 2011.
- 36 Mihai Pătraşcu, L. Roditty, and M. Thorup. A new infinity of distance oracles for sparse graphs. In *Proceedings of Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 738–747, 2012.
- 37 Mihai Pătraşcu and Liam Roditty. Distance oracles beyond the thorup–zwick bound. *SIAM Journal on Computing*, 43(1):300–311, 2014.
- 38 M. Sohel Rahman and Costas S. Iliopoulos. Pattern matching algorithms with don’t cares. In *SOFSEM 2007: Theory and Practice of Computer Science, 33rd Conference on Current Trends in Theory and Practice of Computer Science, Harrachov, Czech Republic, January 20-26, 2007, Proceedings Volume II*, pages 116–126, 2007.
- 39 Robert Endre Tarjan. A class of algorithms which require nonlinear time to maintain disjoint sets. *Journal of Computer and System Sciences*, 18(2):110–127, 1979.
- 40 Hing-Fung Ting and Yilin Yang. Dictionary matching with uneven gaps. In *Annual Symposium on Combinatorial Pattern Matching (CPM)*, volume 9133, page 247, 2015.