

ALGORITHMIC COMPLEXITY OF PROTEIN IDENTIFICATION: SEARCHING IN WEIGHTED STRINGS

Mark Cieliebak, Zsuzsanna Lipták, Emo Welzl
ETH Zurich, Institute of Theoretical Computer Science
cielieba,zsuzsa,emo@inf.ethz.ch

Thomas Erlebach
ETH Zurich, Computer Engineering and Networks Laboratory
erlebach@tik.ee.ethz.ch

Jens Stoye*
Max Planck Institute of Molecular Genetics, and Konrad-Zuse-Zentrum (ZIB), Berlin
stoye@techfak.uni-bielefeld.de

Abstract We investigate a problem which arises in computational biology: Given a constant-size alphabet \mathcal{A} with a weight function $\mu : \mathcal{A} \rightarrow \mathbb{N}$, find an efficient data structure and query algorithm solving the following problem: For a string σ over \mathcal{A} and a weight $M \in \mathbb{N}$, decide whether σ contains a substring with weight M (ONE-STRING MASS FINDING PROBLEM). If the answer is **yes**, then we may in addition require a witness, i.e., indices $i \leq j$ such that the substring beginning at position i and ending at position j has weight M . We allow preprocessing of the string, and measure efficiency in two parameters: storage space required for the preprocessed data, and running time of the query algorithm for given M . We are interested in data structures and algorithms requiring subquadratic storage space and sublinear query time, where we measure the input size as the length of the input string. Among others, we present two non-trivial efficient algorithms: LOOKUP solves the problem with $O(n)$ space and $O(\frac{n}{\log n} \cdot \log \log n)$ time; INTERVAL solves the problem for binary alphabets with $O(n)$ storage space in $O(\log n)$ query time. Finally, we introduce other variants of the problem and sketch how our algorithms may be extended for these variants.

Keywords: Weighted Strings, Protein Identification, Database Searching

*Present address: Bielefeld University, Faculty of Technology

1. Introduction

In the present paper, we introduce a combinatorial problem which originates from computational biology: Given a string σ over a weighted alphabet \mathcal{A} , find a data structure and a query algorithm which, for a given weight $M \in \mathbb{N}$, decides whether σ has a (contiguous) substring of weight M . If the answer is **yes**, we may in addition ask for a *witness*, i.e., two positions within σ where a substring with weight M begins and ends. The actual problem in computational biology is to find several masses M_1, \dots, M_m in a database of strings. We concentrate on the one-string problem because algorithms can be easily extended to the multiple-string problem. We formally define the other problem variants at the end of the paper and sketch how extensions may be designed. There are two simple algorithms which solve the one-string problem: One uses linear time for a query and no additional storage space; the other has logarithmic query time, but requires a preprocessing step and additional storage space for the resulting data structure which may be quadratic. Hereby, space and time complexities are measured in the length of the string. We are interested in algorithms that are better than these two: i.e., we allow preprocessing and look for an algorithm where the data structure needs subquadratic space *and* the running time for a query is sublinear.

Formulated in this way, the problem is a purely combinatorial and algorithmic problem: Are there algorithms which allow searching in weighted strings of size n with $o(n^2)$ additional storage space and $o(n)$ query time? If so, can we find a tradeoff between space and time?

The problem differs from traditional string searching problems in one important aspect: While those look for substructures of strings (substrings, non-contiguous subsequences, particular types of substrings such as repeats, palindromes etc.), we are interested only in *weights* of substrings. This means that, on the one hand, we lose a lot of the structure of strings: e.g. the weight of a string is invariant under permutation of letters; on the other, we gain the additional structure of the weight function, such as its additivity. For instance, the problem of searching in $X + Y$, where X and Y are two sets of numbers, turns out to be closely related to our problem (see [Fre75] and [HPSS75]). However, we have been able to extend negative results which have been reached for that problem ([CDF90]): We can show that this approach (using the naïve solution without preprocessing) cannot lead to an efficient algorithm for our problem. We have been unable to find any treatment of our problem in the vast amount of literature on strings (e.g. [AG95, Gus97, Lot97, RS97, CR94]).

Our problem is positioned between the areas of string algorithms, search algorithms, and algebra. We believe that it is not only relevant for computational biology, but that it is also of theoretical interest to the field of combinatorial searching. As far as we are aware, no efficient algorithms have so far been presented for this problem.

We would like to stress at this point that, even though the source of our problem is a biological question, the results we present here are primarily of theoretical interest. The reason is twofold: First, none of the algorithms we

present are efficiently applicable in their current form. LOOKUP requires sublinear query time, but this is a mainly asymptotic result, since the query time only improves for very long input strings. Algorithm INTERVAL, on the other hand, is very efficient both in query time and storage space, but it only works for alphabets of size 2, a case which does not occur in the usual biological setting. The second reason is that all biological data are prone to errors; in fact, there is no such thing as error-free data. Thus, all applications in computational biology need to be highly fault tolerant. Our algorithms can be straightforwardly adapted to become tolerant to measurement errors. However, this aspect is not included in this paper.

Thus, the present paper demonstrates that efficient algorithms for the problem presented are possible; it remains a great challenge to actually find algorithms that are also of practical value.

1.1. Biological Motivation

Proteomics is the field that investigates the nature and function of proteins. As in molecular biology in general, large amounts of data are being accumulated at present, which presents particular computational and mathematical challenges.

Proteins are large molecules that play a fundamental role in all living organisms. They are made up of smaller molecules (amino acids) that are linked together in a certain order. The sequence of amino acids constitutes the so-called *primary structure* of a protein. Protein size ranges from below 100 to several thousand amino acids, where a typical protein has length 300 to 600. Most proteins are made up of the 20 most common amino acids. For the purposes of this paper, we will view a protein as a finite string over an alphabet of size 20.

The information about known proteins is stored in large databases, such as SWISS-PROT (nearly 100,000 proteins) or PIR (more than 200,000 proteins). When a protein is isolated, one would like to know whether it is already known and if so, to identify it. An obvious way is to establish its primary structure: This is called *de novo protein sequencing*. However, protein sequencing, unlike DNA sequencing, is very expensive (both in time and money!). E.g. identifying *one* amino acid with Edman Degradation, one standard method for protein sequencing, takes about 45 minutes, which makes this approach unfeasible in a high-throughput context.

Therefore, methods are required that test the protein against a database without having to sequence it first. One such method—which we will investigate here—makes use of the differences in molecular weights of amino acids: The protein is broken up into smaller pieces and these pieces are then weighed¹, using a mass spectrometer. This will yield a “fingerprint” of the protein that

¹Biologists will excuse some rough simplifications.

can then be tested against the database: The goal is to find a protein in the database that has substrings matching each of the input masses.

The method used for breaking up the protein into smaller pieces is referred to by biologists as *digestion*: a so-called cleavage agent such as an enzyme, e.g. trypsin, is used which literally cuts the protein in certain well-defined places. The whole process is called *mass spectrometry*. Using digestion is algorithmically rather simple, at least with error-free data, since the breakup points are known in advance; it is thus possible to preprocess the database in an appropriate way. The complications arise due to measurement errors and post-translational modifications. There is a large amount of literature on mass spectrometry [YISGH93, HBS⁺93, JQCG93, PHB93, MHR93, EMYI94]; some papers dealing with different aspects and modifications of the problem, e.g. the minimum number of masses needed to identify a protein [PHB93], combinatorial [PDT00] or probabilistic [BE01] models for scoring the difference of two mass spectra, or approaches for a correct identification even in the presence of post-translational modifications of the protein [MW94, YEM95, PMDT01]. The review [YI98] as well as chapter 11 of the book [Pev00] contain more detailed introductions to this topic. For an introduction to computational biology in general, see [SM97]; for more on molecular biology, [Str88]; while [GW91] is an easy-going introduction to genetics for non-biologists.

In this paper, we deal with algorithmic questions that arise if nothing is known about the breaking points, i.e., we assume random fragmentation. Testing for random weights is algorithmically far more complex than the digestion method, because the cutting places are not known in advance. This approach makes sense because it allows combination of several cleavage agents and it eliminates problems such as incomplete digestion. In addition, since we never make any assumptions about the probability distribution of breaking points, any algorithm for the random fragmentation method can be used for digested inputs, too. In the long run, however, for the biological application, algorithms are needed that are not only efficient, but also fault tolerant: They need to be tolerant both to measurement errors ($M \pm \epsilon$; missing or additional masses in the spectrum) and to sequencing errors.

1.2. Overview

The paper is organized as follows. We first introduce the problem and all necessary definitions in Section 2, where we also present some simple ideas that motivate our efficiency requirements. In Section 3, we design an algorithm (LOOKUP) that is asymptotically efficient, with linear storage space and a query running time that is only just sublinear. LOOKUP thus serves to demonstrate that the requirements we defined earlier can be met. Section 4 contains an algorithm (INTERVAL) that solves the problem for alphabets of size 2 and has a very good performance. However, we do not think that it can be generalized to larger alphabets. In Section 5, we present two other problem variants and discuss how algorithms for the original problem can be extended to these. In

addition, we sketch improvements of our algorithms for special cases. Finally, in Section 6, we sketch ongoing work.

2. Problem and Simple Solutions

Fix an *alphabet* \mathcal{A} of size $|\mathcal{A}| = s$ and a *mass function* $\mu : \mathcal{A} \rightarrow \mathbb{N}$. The *mass* (or the *weight*) of a string (or a *sequence*) σ over \mathcal{A} is defined as the sum of the individual masses $\mu(\sigma) := \sum_{i=1}^n \mu(\sigma(i))$, where $\sigma(i)$ denotes the i 'th letter of σ , and $n = |\sigma|$ is the length of σ . For a mass $M \in \mathbb{N}$ and a string σ of length n , we say that M is a *submass* of σ if σ has a substring of mass M , i.e., if there are indices $1 \leq i \leq j \leq n$ s.t. $\mu(\sigma(i, j)) = M$, where $\sigma(i, j)$ is the substring of σ starting in position i and ending in position j . For $a \in \mathcal{A}$, let us denote the *multiplicity* of a in σ by $|\sigma|_a := |\{i \mid \sigma(i) = a\}|$.

The ONE-STRING MASS FINDING PROBLEM is defined as follows:

Given a string σ of length $|\sigma| = n$ and a mass M , is M a submass of σ ?

A simple algorithm to solve the problem is LINSEARCH, which performs a linear search through the string: For given σ , start at position $\sigma(1)$ and add up masses until reaching the first position j s.t. $\mu(\sigma(1, j)) \geq M$. If the mass of the substring $\sigma(1, j)$ equals M , then output **yes** and stop; else start subtracting masses from the beginning of the string until the smallest index i s.t. $\mu(\sigma(i, j)) \leq M$ is reached. Repeat until finding a pair of indices (i, j) s.t. $\mu(\sigma(i, j)) = M$, or until reaching the end of the string (i.e., until the current substring is $\sigma(i, n)$ for some i and $\mu(\sigma(i, n)) < M$). The algorithm can be visualized as shifting two pointers ℓ and r through the string, where ℓ points to the beginning of the current substring and r to its end. LINSEARCH takes $O(n)$ time, since it looks at each letter at most twice. If we do not allow any preprocessing, this is asymptotically optimal, since it may be necessary to look at each letter at least once.

On the other hand, if preprocessing of σ is allowed, then there is another simple algorithm for the ONE-STRING MASS FINDING PROBLEM which uses binary search: in a preprocessing step, it calculates the set of all possible submasses of σ (i.e., $\mu(\sigma(i, j))$ for all $1 \leq i \leq j \leq n$) and stores them in a sorted array. Given a query mass M , it performs binary search for M in this array. We will refer to this algorithm as BINSEARCH. The space required to store the sorted array is proportional to the number of different submasses in σ , which is bounded by $O(n^2)$. The time for answering a query is thus $O(\log n)$.

From now on, an algorithm for the ONE-STRING MASS FINDING PROBLEM will consist of three components: a preprocessing phase, a data structure in which the result of the preprocessing is stored, and a query method. For a string σ , the preprocessing will be done only once, while the query step will typically be repeated many times. For this reason, we are interested in algorithms with fast query methods, whereas we ignore time and space required for the preprocessing step (as long as they are within reasonable bounds). Space efficiency is measured in storage space required by the data structure.

We are looking for algorithms that are better than LINSEARCH and BINSEARCH, i.e., require storage space $o(n^2)$ for the data structure, and query time $o(n)$. We will call an algorithm *skinny* if the associated data structure requires $o(n^2)$ space, and *speedy* if the query method runs in time $o(n)$.

In this context, the question naturally arises whether a given mass M can be the weight of a string. If the size of the alphabet is variable, then this question is a variant of the INTEGER KNAPSACK PROBLEM, and is NP-complete (cf. [GJ79]). If the alphabet size is constant, the question can be solved with a simple Integer Linear Program.

A third simple algorithm for the ONE-STRING MASS FINDING PROBLEM, which we will call BOOLEANARRAY, works as follows: In the preprocessing phase, define $W := \max\{\mu(a) \mid a \in \mathcal{A}\}$ and let B be a Boolean array of length $n \cdot W$. Set $B[k]$ to **true** if and only if k is a submass of σ . Given a query mass M , we output $B[M]$. This algorithm has query time $O(1)$, while the data structure B requires space $O(n \cdot W)$ bits. Thus, the algorithm is speedy and, if $W = o(n)$, it is skinny, too. However, this does not solve the ONE-STRING MASS FINDING PROBLEM in general, since we do not want to restrict the size of W .

In the following, we assume that the alphabet \mathcal{A} is of constant size and we do not restrict the maximum weight W of a letter. We assume a machine model with word size $\Omega(\log n + \log W)$ in which arithmetic operations on numbers with $O(\log n + \log W)$ bits can be executed in constant time; storage space is measured in terms of the number of words used. Without this assumption, we would get an extra factor $O(\log W)$ in the query time and in the storage space.

3. An Algorithm that is Both Skinny and Speedy

In this section, we present algorithm LOOKUP that solves the ONE-STRING MASS FINDING PROBLEM with storage space $O(n)$ and query time $O(\frac{n}{\log n} \cdot \log \log n)$. The idea is as follows. Similar to the simple linear search algorithm LINSEARCH introduced in Section 2, LOOKUP shifts two pointers along the sequence which point to the potential beginning and end of a substring with mass M . However, $c(n)$ steps of the simple algorithm are bundled into one step here. If $c(n)$ is chosen appropriately, i.e. approximately $\log n$, then this will reduce the number of steps from $O(n)$ to $O(\frac{n}{\log n})$, while each step will now require $O(\log \log n)$ time instead of constant time. We will hereby heavily exploit the fact that the alphabet has constant size.

3.1. An Example

Example 1. Let $\mathcal{A} = \{a, b, c\}$, $\mu(a) = 1$, $\mu(b) = 2$, $\mu(c) = 5$. Let us assume that we are looking for $M = 14$ in $\sigma = abbcabccaabb$. LINSEARCH would shift two pointers ℓ and r through the sequence until reaching positions 5 and 9 respectively. Here, it would stop because the substring $\sigma(5, 9) = abcca$ has weight 14. Let us assume that $c(n) = 3$. We divide the sequence σ into blocks

of size $c(n)$. Now, rather than shifting the two pointers letter by letter, we will shift them by a complete block at a time. In order to do this, for each block we store a pointer to an index I which corresponds to the substring which starts with the first letter of the block and ends with the last. Let us assume now that ℓ is at the beginning of the first block, and r is at the end of the second block, as indicated in Figure 1. We are interested in the possible *changes* to the current submass if we shift the two pointers at most $c(n)$ to the right. Given a list of these, we could search for $M - \mu(\sigma(\ell, r))$. For example, the current submass in Figure 1 is $\mu(\sigma(1, 6)) = 13$, and we want to know whether, by moving ℓ and r at most 3 positions to the right, we can achieve a change of $14 - 13 = 1$.




Figure 1. Example 1 – LOOKUP searching for $M = 14$

We can calculate these possible changes and store them in a $(c(n) + 1) \times (c(n) + 1)$ matrix T whose (i, j) -entry holds the submass change when ℓ is moved $i - 1$ positions to the right, and r is moved $j - 1$ positions to the right:

$$T[abb, cca] = \begin{pmatrix} 0 & 5 & 10 & 11 \\ -1 & 4 & 9 & 10 \\ -3 & 2 & 7 & 8 \\ -5 & 0 & 5 & 6 \end{pmatrix}$$

In order to be able to do fast search, we store the entries of the matrix in a sorted array: $S[abb, cca] = [-5, -3, -1, 0, 2, 4, 5, 6, 7, 8, 9, 10, 11]$. Now we can find out in time $O(\log(\text{size of array}))$ whether the difference we are looking for is there. In the present case, 1 is not in the array, which tells us that we have to move one of the two pointers to the next block.

To determine which pointer to move, we consider what the linear search algorithm LINSEARCH would do when searching for M and starting in the current positions of the left resp. right pointer. Since M is not present within these two blocks, at least one of the two pointers would reach the end of its current block. Here, we want to move the pointer which would first reach the end of its block. We can determine which pointer this is if we compare the difference $M - \mu(\sigma(\ell, r))$ with the matrix entry $T(c(n), c(n))$ corresponding to $c(n) - 1$ moves of both the left and the right pointer (in this case 7). If the difference is smaller, we move the left pointer to the next block, otherwise we move the right one. In our example, we have a difference of 1, thus we move the left pointer to the next block.

This will change the current submass by -5 (the minimum of the array), yielding $\mu(\sigma(4, 6)) = 13 - 5 = 8$. Thus, we now look for $M - \mu(\sigma(4, 6)) =$

$14 - 8 = 6$. The sorted array for this pair of positions is $S[cab, cca] = [-8, -6, -5, -3, -1, 0, 2, 3, 4, 5, 6, 10, 11]$, and the matrix is as follows:

$$T[cab, cca] = \begin{pmatrix} 0 & 5 & 10 & 11 \\ -5 & 0 & 5 & 6 \\ -6 & -1 & 4 & 5 \\ -8 & -3 & 2 & 3 \end{pmatrix}$$

Value 6 is in the array: By looking in the matrix, we can see that a difference of 6 can be achieved by moving the left pointer by 1 position and the right pointer by 3 positions. The algorithm outputs positions 5 and 9 and then terminates.

3.2. Algorithm LOOKUP

We postpone the exact choice of the function $c(n)$ to the analysis, but assume for now that it is approximately $\log n$. For simplicity, we assume that $c(n)$ is a divisor of n .

Preprocessing: Given σ of length n , first compute $c(n)$. Next, build a table T of size $|\mathcal{A}|^{c(n)} \times |\mathcal{A}|^{c(n)}$. Each row resp. column of T will be indexed by a string from $\mathcal{A}^{c(n)}$. For $I, J \in \mathcal{A}^{c(n)}$, the table entry $T[I, J]$ contains the matrix and the sorted array as described above. The matrix contains all differences $\mu(\text{prefix}(J)) - \mu(\text{prefix}(I))$. Note that the table T depends only on n and \mathcal{A} , and not on the sequence σ itself. Next, divide σ into blocks of length $c(n)$. For each block, store a pointer to an index I that will be used to look up table T . Each such index I represents one string from $\mathcal{A}^{c(n)}$.

Query Algorithm: Given M , set $\ell := 1$ and $r := 0$. Repeat the following steps until M has been found or $r > n$:

Step 1: Say ℓ is set to the beginning of the i 'th block and r to the end of the $(j - 1)$ 'th block. Then look in the sorted array in $T(I, J)$ where the pointer of block i resp. j points to index I resp. J . Find whether $M - \mu(\sigma(\ell, r))$ is in the array with binary search.

Step 2: If $M - \mu(\sigma(\ell, r))$ is in the array, search for an entry (u, v) in the matrix $T(I, J)$ which equals $M - \mu(\sigma(\ell, r))$ by exhaustive search², and return **yes**, along with the witness $i' := (i - 1) \cdot c(n) + u, j' := (j - 1) \cdot c(n) + (v - 1)$, since $\mu(\sigma(i', j'))$ has mass M .

Step 3: Otherwise, $M - \mu(\sigma(\ell, r))$ is not in the array. If $M - \mu(\sigma(\ell, r))$ is less than the matrix entry at position $(c(n), c(n))$, then increment ℓ by $c(n)$ and set $\mu(\sigma(\ell, r)) := \mu(\sigma(\ell, r)) + \min(\text{array})$; otherwise, increment r by $c(n)$ and set $\mu(\sigma(\ell, r)) := \mu(\sigma(\ell, r)) + \max(\text{array})$.

Analysis: First we derive formulas for space and time, and then we show how to choose $c(n)$. The space needed for storing table T is $|\mathcal{A}|^{2c(n)} \cdot ((c(n) + 1)^2 + (c(n) + 1)^2) = O(|\mathcal{A}|^{2c(n)} \cdot c(n)^2)$. Space needed for storing the pointer at each block is $\frac{n}{c(n)} \cdot \log(|\mathcal{A}|^{c(n)}) = O(n)$. For the last equality, recall that \mathcal{A} is of constant size. For the query time, observe that after each iteration

²Alternatively, we could have stored (u, v) during the preprocessing in the sorted array, too.

(consisting of Steps 1 to 3), either ℓ or r is advanced to the next block. As each of the pointers can advance at most $\frac{n}{c(n)}$ times, there can be at most $2\frac{n}{c(n)}$ iterations. Each iteration except the last one takes time $O(\log c(n)^2) + O(1)$. The last iteration may take time $O(c(n)^2)$.

In total, the algorithm requires storage space $O(n + |\mathcal{A}|^{2c(n)} \cdot c(n)^2)$ and time $O(\frac{n}{c(n)} + \frac{n}{c(n)} \log c(n) + c(n)^2)$. Now, if we choose $c(n) := \frac{\log_{|\mathcal{A}|} n}{4}$, then we obtain $|\mathcal{A}|^{c(n)} = n^{\frac{1}{4}}$. This yields a storage space of $O(n + n^{\frac{1}{2}} \cdot \log^2 n) = O(n)$ and query time $O(\frac{n}{\log n} \log \log n)$, which is both skinny and speedy. Other choices of $c(n)$ do not asymptotically improve time and space at the same time.

Theorem 1. *Algorithm LOOKUP solves the ONE-STRING MASS FINDING PROBLEM with storage space $O(n)$ and query time $O(\frac{n}{\log n} \log \log n)$.*

Thus, LOOKUP beats both the query time of LINSEARCH and the storage space of BINSEARCH. However, its practical use is limited to very long sequences: In order to obtain a block size of, say, $c(n) = 10$, the input string would have to have length $n = |\mathcal{A}|^{40}$. In the next section we present a practical algorithm for binary alphabets.

4. A Speedier Algorithm for Binary Alphabets


In this section, we present algorithm INTERVAL which solves the ONE-STRING MASS FINDING PROBLEM for an alphabet of size 2. It uses storage space $O(n)$ and has query time $O(\log n)$. The algorithm *decides* whether a given mass is a submass of σ , but does not return a witness.

Let σ be a string over $\mathcal{A} := \{a, b\}$ of length n and fix $k \leq n$. Observe that, when sliding a window of size k over σ , then in one step, the multiplicities of a and b within the window change at most by one. We represent substrings of σ by points in the $\mathbb{Z} \times \mathbb{Z}$ lattice, where the two coordinates signify the multiplicities of a and b :

$$S_k := \{(i, j) \in \mathbb{Z} \times \mathbb{Z} \mid i + j = k, \text{ there is a substring } \tau \text{ of } \sigma : |\tau|_a = i, |\tau|_b = j\}.$$

All points in S_k will lie on a line (a diagonal), and moreover, they will be neighbours. We will refer to such a set of neighbours on a line as an *interval*. Each such interval has two extremal points.

Example 2. $\sigma = aaaaabaabb$. The figure shows the representation of all substrings of length $k = 8$. Extremal points of this interval are $(5, 3)$ and $(7, 1)$.



Assume for a moment that we know the multiplicities of a and b in M , e.g. $M = i \cdot \mu(a) + j \cdot \mu(b)$. Then we can easily find out whether M is a submass of σ : We store the S_k 's, for $1 \leq k \leq n$ by their extremal points during the

preprocessing phase. Now we only have to check whether $(i, j) \in S_{i+j}$, which takes $O(1)$ time. This requires storage space linear in n . If, in addition, i and j were known to be the only feasible multiplicities of a and b (i.e., the unique solution of the equation $x \cdot \mu(a) + y \cdot \mu(b) = M$), then this algorithm would even decide whether M is a submass of σ , and we would be done.

Unfortunately, we do not know the multiplicities of a and b in M . We define $d := \mu(b) - \mu(a)$ (w.l.o.g., assume $\mu(a) < \mu(b)$) and use the residue of $M \bmod d$ to look up a table. The table, generated during the preprocessing phase, contains representations of all submasses of σ .

Let $M_k := \{\mu(\tau) \mid \tau \text{ is a } k\text{-length substring of } \sigma\}$. Observe that consecutive elements of M_k (when sorted) differ by exactly d . Therefore, we can write $M_k = \{c_k + \ell \cdot d \mid \ell = 0, \dots, n_k - 1\}$, where $c_k = \min M_k$ and $n_k = |M_k|$. Furthermore, $M_k = \{r_k + \ell \cdot d \mid \ell = a_k, \dots, b_k\}$, where $r_k := (c_k \bmod d)$, $a_k := \lfloor \frac{c_k}{d} \rfloor$ and $b_k := a_k + n_k - 1$. This says that all submasses of the same length have the same residue modulo d .

Observe that $r_k = (k \cdot \mu(a) \bmod d)$. Thus, we may have the same residue modulo d for different values of k . Instead of storing a_k and b_k for each r_k individually (which could result in linear query time), we will store the union of all intervals which belong to the same residue r , sorted by their endpoints. For an example, see [CEL⁺01].

4.1. Algorithm INTERVAL

In the preprocessing phase, we calculate the r_k 's, a_k 's, and b_k 's as above. We then sort the r_k 's, thus obtaining a sorted array q_1, \dots, q_m , where $m \leq n$ (since different S_k 's may have the same residue). For each q_l , we compute a list of interval endpoints which represents the union of all intervals $[a_k, b_k]$ with $r_k = q_l$. This list consists of one or more disjoint intervals, which we store in sorted order in an array A_l .

Now, when querying whether a given mass M is contained in σ , first decompose $M = g \cdot d + r$, where $r = (M \bmod d)$ and $g \in \mathbb{N}$; then find index $l \in \{1, \dots, m\}$ such that $r = q_l$, using binary search; if no such index can be found, then M is not a submass of σ , and the algorithm outputs **no**; otherwise, find whether there is an interval $[a, b]$ in array A_l such that $g \in [a, b]$, using binary search on (the left endpoints of) the intervals; M is a submass of σ if and only if such an interval exists.

Since the total number of intervals to be stored is n , the storage space needed is $O(n)$. The first step of the query algorithm takes time $O(1)$. The second step takes time $O(\log n)$, since the number of different residues is at most n . The third step takes time $O(\log n)$, since the maximum number of intervals stored in one array A_l is n . We obtain a total query time $O(\log n)$.

Theorem 2. *Algorithm INTERVAL solves the ONE-STRING MASS FINDING PROBLEM for binary alphabets with storage space $O(n)$ and query time $O(\log n)$.*

The problem in generalizing this approach to larger alphabets is that the algorithm relies on the crucial fact that points representing substrings of the

same length lie on a line and form an interval. This does not generalize to higher dimensions, since there we only know that the points representing substrings of the same length are connected.

5. Problem Variants

The MULTIPLE-STRING MASS FINDING PROBLEM is defined as follows:

Given k strings $\sigma_1, \dots, \sigma_k$ and a mass $M \in \mathbb{N}$, return a list i_1, \dots, i_r of those strings σ_{i_j} which have M as a submass.

An algorithm Ψ for the ONE-STRING MASS FINDING PROBLEM can be extended to an algorithm for the MULTIPLE-STRING MASS FINDING PROBLEM by running Ψ on each string σ_i one by one. Required storage space and query time simply sum up.

Alternatively, we can adapt an approach from Group Testing (cf. [DH00]): We define a new string $\sigma := \sigma_1\omega\sigma_2\omega\dots\omega\sigma_k$, where ω is a new letter with mass $\mu(\omega) := \max\{\mu(\sigma_i) \mid 1 \leq i \leq k\} + 1$. Before applying Ψ to σ , we check whether $M \geq \mu(\omega)$. If so, then M cannot be a submass of any of the strings, and we are done. Otherwise, we know that whenever Ψ finds mass M in σ , then it is a submass of σ_i for some index i . If algorithm Ψ can output *all* positions of M in σ , this solves the MULTIPLE-STRING MASS FINDING PROBLEM. If Ψ only *decides* whether M is a submass of σ (i.e. it outputs only **yes** or **no**), we use a kind of “binary tree search” BINTREESEARCH to find all σ_i with submass M as follows. First, we run Ψ on σ as described above. If it outputs **no**, then no string σ_i has submass M , and we are done. Otherwise, we divide σ into two new strings $\sigma_l := \sigma_1\omega\dots\omega\sigma_{\lfloor \frac{k}{2} \rfloor}$ and $\sigma_r := \sigma_{\lfloor \frac{k}{2} \rfloor + 1}\omega\dots\omega\sigma_k$ and run Ψ on both strings separately. We repeat the division step until the new strings cover exactly one σ_i , in which case the answer of Ψ determines whether σ_i has a submass M . Analysis of BINTREESEARCH depends heavily on storage space and query time required by Ψ . For instance, if algorithm Ψ requires storage space linear in the length of the string, then the storage space of BINTREESEARCH is $O((\log k) \cdot \sum_{i=1}^k |\sigma_i|)$. Query time of BINTREESEARCH depends on the number of strings with submass M , in contrast to the simple idea of applying Ψ to each string separately.

Given a specific algorithm for the ONE-STRING MASS FINDING PROBLEM, there might be even better ways to extend it to the MULTIPLE-STRING MASS FINDING PROBLEM: E.g. for BINSEARCH, we can use *one* sorted array to store all submasses of all strings. For each submass x we store the set of indices I_x of all those strings which have a submass x . Given mass M , we perform binary search in the array and output all indices stored in I_M . Required storage space remains unchanged, but the running time becomes $O(\log(\sum_{i=1}^k |\sigma_i|) + |I_M|)$, where $|I_M| \leq k$ is the size of the output. A similar idea applies to LOOKUP, where we could store only one table T of size $|\mathcal{A}|^{c_{max}} \times |\mathcal{A}|^{c_{max}}$ where $c_{max} = \max_{i=1}^k c(|\sigma_i|)$, and use it for all runs of the algorithm. However, this does not decrease the asymptotic space required, which still remains linear.

We define a third problem variant, the **MULTIPLE-STRING MULTIPLE-MASS FINDING PROBLEM**:

Given k strings $\sigma_1, \dots, \sigma_k$, m masses $M_1, \dots, M_m \in \mathbb{N}$, and a threshold $1 \leq t \leq m$, return a list i_1, \dots, i_r of those strings σ_{i_j} which have at least t of the masses as submasses.

In the setting of our application in computational biology, this will be a more realistic formulation, since typically, one breaks a given protein in several pieces and wants to find the protein in the database which contains all (or at least many) of these pieces. Obviously, the **MULTIPLE-STRING MULTIPLE-MASS FINDING PROBLEM** can be solved by applying algorithms for the **MULTIPLE-STRING MASS FINDING PROBLEM** m times. We are investigating the question whether concurrently searching for m masses can be performed more efficiently.

Finally, we present an improvement of all our algorithms for “short masses”: Let the *length* of a mass M be defined as $\lambda(M) := \max(\{|\tau| \mid \tau \in \mathcal{A}^*, \mu(\tau) = M\} \cup \{-1\})$. Here, $\lambda(M) = -1$ means that there is no string with mass M . Suppose that we know in advance that all query masses are short in comparison to n , i.e., that there is a function $f(n)$ such that $\lambda(M) \leq f(n) = o(n)$ for all queries M . Then there is a simple algorithm to solve the **ONE-STRING MASS FINDING PROBLEM**, which is a variant of **BINSEARCH**: In the preprocessing, we store all submasses of σ of length $\ell \leq f(n)$ in a sorted array. This requires storage space $O(n \cdot f(n))$, since for each position i in σ , at most $f(n)$ substrings of length $\ell \leq f(n)$ start in i . For a query, we do binary search in this array. This takes time $O(\log n)$, which is speedy. Since $f(n) = o(n)$, the algorithm is skinny, too. We can use this approach to improve our algorithms in the sense that they will run faster on short masses.

6. Conclusion

For some algorithms (e.g. **BINSEARCH**), storage space and query time for a string σ depend on the number of different submasses of σ . This number is bounded from below by $\Omega(n)$ and from above by $O(n^2)$, and there are examples which meet these boundaries: For instance, let $\mathcal{A} = \{a, b\}$, $\mu(a) = 1$ and $\mu(b) = n + 1$. The string a^n has n different submasses, while the string $a^n b^n$ has $(n + 1)^2 - 1$ different submasses. It may be interesting to explore properties of the number of different submasses, e.g. its expected value or a characterization of all strings with $\Theta(n^2)$ submasses. Furthermore, we are looking for efficient algorithms to compute this number.

With **LOOKUP**, we presented an algorithm for the **ONE-STRING MASS FINDING PROBLEM** which is both skinny and speedy. This proves that it is asymptotically possible to beat both **LINSEARCH** and **BINSEARCH** at the same time. This raises the question whether there are more practical algorithms that are skinny and speedy. In the long run, we are interested in the tradeoff between query time and storage space for the **ONE-STRING MASS FINDING PROBLEM**. Do algorithms exist that can be parametrized to allow for adjustment of this tradeoff?

Acknowledgments

We would like to thank Juraj Hromkovic who read an earlier version of this paper and made many helpful suggestions, and Sacha Baginsky for comments on the biological introduction.

References

- [AG95] A. Apostolico and Z. Galil, editors. *Combinatorial Algorithms on Words*. Springer, 1995.
- [BE01] V. Bafna and N. Edwards. SCOPE: A probabilistic model for scoring tandem mass spectra against a peptide database. *Bioinformatics*, 17(Supplement 1):S13–S21, 2001.
- [CDF90] M. Cosnard, J. Duprat, and A. G. Ferreira. The complexity of searching in $X + Y$ and other multisets. *Information Processing Letters*, 34:103–109, 1990.
- [CEL⁺01] M. Cieliebak, T. Erlebach, Zs. Lipták, J. Stoye, and E. Welzl. Algorithmic complexity of protein identification: Combinatorics of weighted strings. Technical Report 361, ETH Zurich, Department of Computer Science, 2001.
- [CR94] M. Crochemore and W. Rytter. *Text Algorithms*. Oxford University Press, New York, NY, 1994.
- [DH00] D. Du and F. K. Hwang, editors. *Combinatorial Group Testing and its Applications*. World Scientific, second edition, 2000.
- [EMYI94] J. Eng, A. McCormack, and J. R. Yates III. An approach to correlate tandem mass spectral data of peptides with amino acid sequences in a protein database. *J. Amer. Soc. Mass Spect.*, 5:976–989, 1994.
- [Fre75] M. L. Fredman. Two applications of a probabilistic search technique: Sorting $X + Y$ and building balanced search trees. In *Conference Record of Seventh Annual ACM Symposium on Theory of Computing (STOC)*, pages 240–244, 1975.
- [GJ79] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, 1979.
- [Gus97] D. Gusfield. *Algorithms on Strings, Trees, and Sequences*. Cambridge University Press, 1997.
- [GW91] L. Gonick and M. Wheelis. *The Cartoon Guide to Genetics*. Harper-Perennial, updated edition, 1991.
- [HBS⁺93] W. J. Henzel, T. M. Billeci, J. T. Stults, S. C. Wong, C. Grimley, and C. Watanabe. Identifying proteins from two-dimensional gels by molecular mass searching of peptide fragments in protein sequence databases. *Proc. Natl. Acad. Sci. USA*, 90(11):5011–5015, 1993.
- [HPSS75] L.H. Harper, T.H. Payne, J.E. Savage, and E. Straus. Sorting $X + Y$. *Communications of the ACM*, 18(6):347–349, 1975.
- [JQCG93] P. James, M. Quadroni, E. Carafoli, and G. Gonnet. Protein identification by mass profile fingerprinting. *Biochem. Biophys. Res. Commun.*, 195(1):58–64, 1993.

- [Lot97] M. Lothaire. *Combinatorics on Words*. Cambridge University Press, second edition, 1997.
- [MHR93] M. Mann, P. Højrup, and P. Roepstorff. Use of mass spectrometric molecular weight information to identify proteins in sequence databases. *Biol. Mass Spectrom.*, 22(6):338–345, 1993.
- [MW94] M. Mann and M. Wilm. Error-tolerant identification of peptides in sequence databases by peptide sequence tags. *Anal. Chem.*, 66(24):4390–4399, 1994.
- [PDT00] P. A. Pevzner, V. Dančík, and C. L. Tang. Mutation-tolerant protein identification by mass spectrometry. *J. Comp. Biol.*, 7(6):777–787, 2000.
- [Pev00] P. A. Pevzner. *Computational Molecular Biology: An Algorithmic Approach*. MIT Press, 2000.
- [PHB93] D. J. C. Pappin, P. Højrup, and A. J. Bleasby. Rapid identification of proteins by peptide-mass fingerprinting. *Curr. Biol.*, 3(6):327–332, 1993.
- [PMDT01] P. A. Pevzner, Z. Mulyukov, V. Dančík, and C. L. Tang. Efficiency of database search for identification of mutated and modified proteins via mass spectrometry. *Genome Res.*, 11(2):290–299, 2001.
- [RS97] G. Rozenberg and A. Salomaa, editors. *Handbook of Formal Languages*, volume 1-3. Springer, 1997.
- [SM97] J. Setubal and J. Meidanis. *Introduction to Computational Molecular Biology*. PWS Boston, 1997.
- [Str88] L. Stryer. *Biochemistry*. Freeman, 1988.
- [YEM95] J. R. Yates, J. K. Eng, and A. L. McCormack. Mining genomes: Correlating tandem mass-spectra of modified and unmodified peptides to sequences in nucleotide databases. *Anal. Chem.*, 67(18):3202–3210, 1995.
- [YI98] J. R. Yates III. Database searching using mass spectrometry data. *Electrophoresis*, 19(6):893–900, 1998.
- [YISGH93] J. R. Yates III, S. Speicher, P. R. Griffin, and T. Hunkapillar. Peptide mass maps: A highly informative approach to protein identification. *Anal. Biochem.*, 214:397–408, 1993.