| CSC 448 | Bioinformatics Agorithms | Alexander Dekhtyar |
| :--- | :--- | :--- |

String Matching. . .

## String Matching Problems and Bioinformatics

Exact String Matching. Given a string $S=s_{1} s_{2} \ldots s_{n}$ of charachters in some alphabet $\Sigma=\left\{a_{1} \ldots a_{K}\right\}$ and a pattern or query string $P=p_{1} \ldots p_{m}$, $m<n$ in the same alphabet, find all occurrences of $P$ in $S$.
"find all occurrences" $=$ report all positions $i_{1}, \ldots, i_{k}$ such that $s_{i_{j}} \ldots s_{i_{j}+m}=$ $P$.

Approximate String Matching.[4] Given a string $S=s_{1} s_{2} \ldots s_{n}$ of charachters in some alphabet $\Sigma=\left\{a_{1} \ldots a_{K}\right\}$, pattern or query string $P=p_{1} \ldots p_{m}, m<n$ in the same alphabet, a maximum error allowed $k \in \mathbb{R}$, and a distance function $d: \Sigma^{*} \times \Sigma^{*} \rightarrow \mathbb{R}$, find all such positions $i$ in $S$, that for some $j$,

$$
d\left(s_{i} \ldots s_{j}, P\right) \leq k
$$

String matching in bioinformatics. Certain known nucleotide and/or amino acid sequences have properties known to biologists. E.g.,ATG is a string which must be present at the beginning of every protein (gene) a DNA sequence.
A primer is a conserved DNA sequence used in the Polymerase Chain Reaction (PCR) to identify the location of the DNA sequence that will be amplified (amplification starts at the location immediately following the 3 ' end primer, known as the forward primer). Finding if a DNA sequence contains a specific (candidate) primer is therefore paramount to the ability to run correct PCR.

A conserved DNA sequence is a sequence of nucleotides in DNA, which is found in the DNA of multiple species and/or multiple strains (for bacteria/prokaryotes), or, in general, in all/almost all sequences in a specific collection. Some sequences are conserved precisely. However, a lot of sequences are conserved with some modifications. Finding such modified strings is an important process for mapping DNA of a new organism, based on the known DNA of a related organism.

For example, we know that genetic code is redundant. Consider the following sequence of nucleotides:

GCTACTATTTTTCAT
When read in forward frame 0 , this sequence encodes the following sequence of amino acids:

## ATIFH

Because of redundancy of the genetic code, the following sequences of nucleotides will produce the same sequence of amino acids (lower case letters show differences with the original string):

```
GCT ACT ATT TTT CAT
GCc ACc ATT TTa CAc
GCc ACc ATc TTg CAT
GCT ACT ATc TTa CAc
GCg ACc ATa TTc CAc
```

```
A T I F H
```

Therefore, it is often useful to search for approximate string matches in DNA sequences.

## Exact String Matching.

There is a wide range of exact string matching algorithms. We briefly outline the following:

1. Naïve string matching.
2. Rabin-Karp algorithm.
3. Finite-automata-based string matching.
4. Knuth-Morris-Pratt (KMP) algorithm.
5. Boyer-Moore algorithm.
6. Gusfield's $Z$ algorithm.

Algorithm Evaluation. We are interested in the following properties of each algorithm:

- Worst-case algorithmic complexity. The traditional measure of algorithm efficiency.
- Expected algorithmic complexity. Behavior of some algorithms is much better than the worst case on most inputs. (E.g., Rabin-Karp algorithm has the same worst-case complexity as the naïve method), but, in practice, runs much faster on most inputs of interest).
- Space complexity. We prefer our algorithms, even the fast ones to occupy space that is linear in the problem size.

Problem size. The "official" input to all string matching algorithms is a pair of strings $S$ and $P$. Their lenghts, $n$ and $m$, respecitvely, supply the problem size for us.

We note that $m<n$. In some cases, especially in bioinformatics applications $m \ll n$ : our DNA sequences may be millions of nucleotides long, while the pattern $P$ we are looking for (e.g., for primer identification) may be on the order of tens of nucleotides. In such cases, we may consider the size of input to be $n$, rather than $n+m$.

In our studies we fix the alphabet $\Sigma$. In bioinformatics, it makes sense, since the two known alphabets, nucleotide and amino acid ones, have 4 and 20 characters in them repsectively. Occasionally, $\Sigma=K$ may be considered part of the size of input. (in our studies, it won't).

## Naïve Algorithm

The naïve algorithm takes as input strings $S$ and $P$ and checks if $s_{i} \ldots s_{i+m}=P$ for each $i=1, \ldots n-m+1$ directly. The pseudocode for this algorithm looks as follows. Here and in further algorithms we adopt notation $S[x, y]$ to represent the substring $s_{x} \ldots s_{y}$ of $S$.

```
Algorithm NaiveStringMatch(S,P,n,m)
begin
    for \(\mathrm{k}=1\) to \(n-m+1\) do
        if \(P=S[k, k+m-1]\) then
            output(k);
    end for
end
```

Analysis. The if statement in the algorithm takes (up to) $m$ atomic comparisons of the form $P[i] ?=S[k+i-1]$ to perform. The for loop repeats $n-m+1$ times. Therefore the running time of this algorithm is $O((n-m+1) \cdot m)=$ $O\left(n m-m^{2}+m\right)=O(n m)$. Note, that is $m=\frac{n}{2}$, then $O(n m)=O\left(n^{2}\right)$, so the in the worst-case, the naïve algorithm is quadratic in the length of the input string $S$.

## Rabin-Karp Algorithm

Idea. The key ideas behind this algorithm are:

- Comparison of two numbers is up to $m$ times faster than comparison of two strings of length $m$.
- Represent the string $P$ and substrings of $S$ of length $m$ as numbers in base $K$ (size of the alphabet) numeric system. Compare the numbers to check if the strings match.
- If numeric representations of strings start being too large to allow for a single comparison, compare numeric representations of the two strings mod $q$ for some relatively large number $q$.
The latter will require a comparison of the actual strings in case their numeric values co-incide, but in practice there should be very few of such comparisons that would fail.

Numeric representation of strings. Let $\Sigma=\left\{a_{1}, \ldots, a_{K}\right\}$. We associate with each character $a_{i}$ a digit $i-1$ in base $K$ numeric system.

Example. The nucleotide alphabet is $\Sigma=\{A, C, G, T\}$. We can map individual characters in this alphabet to digits in base 4 system as follows:

| Character | Digit |
| :--- | :--- |
| A | 0 |
| C | 1 |
| G | 2 |
| T | 3 |

Consider a string $P=p_{1} \ldots p_{m}$ in alphabet $\Sigma$. Without loss of generality we use $p_{i}$ to refer both to the character in $\Sigma$ and to the digit in the base $K$ alphabet, associated with it. The numeric $p$ value represented by $P$ can be computed using the following expression:

$$
p=p_{m}+K\left(p_{m-1}+K\left(p_{m-2}+\ldots+K\left(p_{2}+K p_{1}\right)\right) \ldots\right)
$$

Given a string $S=s_{1} \ldots s_{n}$, where $n>m$, the numeric value of $S[1, m], t_{0}$ can be computed in the same way:

$$
t_{0}=s_{m}+K\left(s_{m-1}+K\left(s_{m-2}+\ldots+K\left(s_{2}+K s_{1}\right)\right) \ldots\right)
$$

More importantly, given $t_{i-1}$, the numeric value $t_{i}$ representing substring $S[i+1, i+m]$ can be computed incrementally as follows:

$$
t_{i}=K\left(t_{i-1}-K^{m-1} s_{i}\right)+s_{i+m}
$$

We can precompute $K^{m-1}$ as this number is a constant.

Example. Consider for following string $S=$ "ATTCCGT". Suppose we are looking for four-letter substrings. There are $|S|-4+1=7-4+1=4$ four-letter substrings in $S: S[1,4], S[2,5], S[3,6]$ and $S[4,7]$. We compute $t_{0}$ for $S[1,4]$ using the mapping of nucleotide symbols to digits in base 4 numeric system as follows:
$S[1,4]=$ "ATTC" $; t_{0}=1+4(3+4(3+4 \cdot 0))=1+4(3+4 \cdot 3)=1+4 \cdot 15=61$.

Further, we compute $t_{1}, t_{2}$ and $t_{3}$ as follows:

$$
\begin{aligned}
& S[2,5]=\text { "TTCC" } ; \quad t_{1}=4\left(t_{0}-4^{3} \cdot 0\right)+1=4(61-0)+1=244+1=245 . \\
& S[3,6]=\text { "TCCG" } ; \quad t_{2}=4\left(t_{1}-4^{3} \cdot 3\right)+2=4(245-64 \cdot 3)+2=4(245-192)+2=4 \cdot 53+2=214 . \\
& S[4,7]=\text { "CCGT" } ; \quad t_{3}=4\left(t_{2}-4^{3} \cdot 3\right)+3=4(214-192)+3=4 \cdot 22+3=91 .
\end{aligned}
$$

Algorithm in a nutshell. Given $P$ and $S$, we compute the numeric values $p$ for $P$ and $t_{0}$ for $S[1, m]$. We compare $p$ to $t_{0}$, and then, using the incremental recompuation method, compute $t_{1}, t_{2}$, etc. and compare them to $p$. Any time $p=t_{i}$ for some $i$, we report $i+1$ as the position of a match.

Complication. This version of Rabin-Karp algorithm works as desired, when $m$ is small enough, that $p$ and $t_{i}$ can fit the machine word. E.g.,

$$
4^{50}=1,267,650,600,228,229,401,496,703,205,376 \sim 10^{31}
$$

This means that strings $P$ up to length 51 can be found in larger strings on processors with 32 -bit architectures. For longer strings, the algorithm will have to be modified.

Modification. Pick a number $q$, and perform computations of $p$ and $t_{i}$ values $\bmod q$. If $p \neq t_{i}(\bmod q)$, then $S[i+1, i+m] \neq P$. However, $p=t_{i}$ does not immediately guarantee $S[i+1, i+m]=P$. We run the full comparison of $P$ and $S[i+1, i+m]$. The pseudocode for the modified algorithm is as follows. The algorithm takes as input, in addition to strings $P$ and $S$ and their respective lenghts $m$ and $n$, two more numbers: $K$, the size of the alphabet and $q$.

```
Algorithm RabinKarp(S, P, n, m, K, q)
begin
    \(h \leftarrow K^{m-1} \bmod q ;\)
    \(p \leftarrow 0 ;\)
    \(t_{0} \leftarrow 0 ;\)
    for \(\mathrm{i}=1\) to \(m\) do //compute initial numeric values
        \(p \leftarrow(K \cdot p+P[i]) \bmod q ;\)
        \(t_{0} \leftarrow\left(K \cdot t_{0}+S[i]\right) \bmod q ;\)
    end for
    for \(\mathrm{j}=0\) to \(n-m\) do //search for substring P in S
        if \(p=t_{j}\) then //if numeric values match,
            if \(P=S[j+1, j+m]\) then need to compare substrings
            output \(j+1\);
            end if
        end if
        if \(j<n-m\) then //move onto the next substring in S
            \(t_{j+1}=\left(\left(K\left(t_{j}-S[j+1] \cdot h\right)+S[j+m+1]\right) \bmod q ;\right.\)
    end for
end
```

Analysis. The if $P=S[j+1, j+m]$ comparison costs $O(m)$ atomic comparisons. In the worst case (e.g., when $S=A^{n}$ and $P=A^{m}$ ), this line will be executed on each iteration of the for loop. There are $n-m+1$ iterations, and therefore, the worst-case complexity of this algorthm is $O((n-m+1) \cdot m)=$ $O(n m)$. (We also note that the first loop has running time $O(m)$, and therefore does not contribute significantly to the overall complexity of the algorithm).

However, unlike the naïve algorithm, Rabin-Karp algorithm improves its running time significantly when there are relatively few, comparing to $n$ matches between $p$ and $t_{0}$. E.g., in many cases when $P$ is NOT a substring of $S$, the running time of the algorithm will be $O(m+n)$ (the first loop takes $O(m)$, the second $-O(n)$ ).

If $S$ has $v$ valid shifts (i.e., $v$ locations in which $P$ occurs), then we can estimate the running time of the Rabin-Karp algorithm as $O(m)+O(n(v+n / q))$, where $n / q$ is our estimate of the probability of a "misfire" (i.e., $p=t_{i}$, when $P \neq S[i+i, i+m])$.

Choosing q. The best value of $q$ is a prime number that is just a bit smaller than $2(B-1)$ where $B$ is the number of bits in the machine word on a given processor architecture. For 16 -bit words ( 16 -bit INTs), take $q=32,749$. For 32 -bit words, you can use $q=2,147,483,629$ or $q=2,147,483,647=2^{31}-1$.

## Finite-automata-based String Matching

Idea. A string $P=p_{1} \ldots p_{m}$ is a finite string, and therefore, is a regular expression in $\Sigma^{*}$. A regular expression matching all strings in $\Sigma^{*}=\left\{a_{1}, \ldots, a_{K}\right\}$ that have $P$ as a substring is

$$
\left(a_{1}\left|a_{2}\right| \ldots \mid a_{K}\right)^{*} p_{1} p_{2} \ldots p_{m}\left(a_{1}\left|a_{2}\right| \ldots \mid a_{K}\right)^{*}
$$

Therefore, we can use finite automata to check if an input string $S$ has $P$ as a substring.

Algorithm structure. The finite-automata-based string matching consists of two parts:

1. Preprocessing. Given a string $P$, a finite state machine (finite automaton) $M_{P}$ for checking for the occurrence of $P$ is constructed.
2. Matching. Given a string $S$ and $M_{P}, M_{P}$ is run on $S$, and it outputs the index $i-|P|$ for each index $i$ of $S$, at which $M_{P}$ reaches a final state.

We notice that these two stages are completely decoupled. $M_{P}$ can be constructed, given $P$ ahead of time. Also, if multiple strings need to be tested for $P$ being a substring (a situation very common in bioinformatics), the Preprocessing step shall only be executed once.

Finite automata. A finite automaton $M$ is a tuple

$$
M=\left\langle Q, q_{0}, A, \Sigma, \delta\right\rangle
$$

where

- $Q=\left\{q_{1}, \ldots, q_{l}\right\}$ is called a set of states.
- $q_{0} \in Q$, is called the start state.
- $A \subseteq Q$ is the set of accepting or final states.
- $\Sigma=\left\{a_{1}, \ldots, a_{K}\right\}$ is a finite input alphabet.
- $\delta: \Sigma \times Q \rightarrow Q$ is the transition function.
$M$ accepts string $S=s_{1} \ldots s_{n}$ iff $\delta\left(s_{n}, \delta\left(s_{n-1}, \delta\left(\ldots \delta\left(s_{1}, q_{0}\right) \ldots\right) \in A\right.\right.$.

Finite Automata for substring matching. We want to design a finite automaton $M_{P}$ for recognizing a given string $P=p_{1} \ldots p_{m}$ as a substring in the input stream. $M_{P}$ is constructed as follows: $M_{P}=\left\langle Q, q_{0}, A, \Sigma, \delta\right\rangle$, where

- $Q=\{0,1, \ldots m\}$.
- $q_{0}=0$.


Figure 1: A finite automaton for matching against the string $P=$ ATAA.

- $A=\{m\}$.
- $\Sigma=\left\{a_{1}, \ldots, a_{K}\right\}$.
- delta is constructed as described below.

Defining the transition function. Clearly, we want $\delta\left(p_{i}, i-1\right)=p_{i}$ for $i=1, \ldots m$.

However, we cannot just say $\delta(a, i-1)=0$ if $a \neq p_{i}$. We may have to move to a different state, because we may still be observing a valid non-empty prefix of $P$.

Example. Consider a string $P=$ ATAA and $S=$ ATATATAA. We note that $p_{1}=$ $s_{1}=A, p_{2}=s_{2}=T$ and $p_{3}=s_{3}=A$. However, $p_{4}=A$, while $s_{4}=T$, so $S[1,4] \neq P:$

$$
\begin{array}{ll}
\text { S: } & \text { ATATATAA } \\
\text { P: } & \text { ATA A }
\end{array}
$$

However, at this point, $s_{3} s_{4}=A T$ form a valid non-empty prefix of $P$, so rather than trying to reapply the entire string $P$ to $S[5,8]$, we must to check if $S[3,8]$ contains $P$, or, knowing that $s_{3} s_{4}=p_{1} p_{2}$ - if $S[5,8]$ contains $p_{3} p_{4}$ as its prefix:

```
S: ATATATAA
P: A T A a
P: AT A a
```

Similar situation occurs on this step as well, and again $s_{5} s_{6}=p_{1} p_{2}$, so we need to check if $S[7,8]$ has $p_{3} p_{4}$ as its prefix:

```
S: A T A T A TAA
P: A T A a
P: A T A a
P: A T A A
```

The finite state machine $M_{P}$ (in the $\{A, T\}$ alphabet for simplicity) is shown in Figure 1.

Constructing the transition function. A suffix function $\sigma: \Sigma^{*} \longrightarrow Q$ is defined as follows:

$$
\begin{aligned}
& \sigma(x)=i \text {, where } p_{1} \ldots p_{i} \text { is the longest suffix of } x \text {, i.e., if } x= \\
& x_{1} \ldots x_{n} \text {, then } \sigma(x)=i \text { means that } x_{n-i} \ldots x_{n}=p_{1} \ldots p_{i} .
\end{aligned}
$$

For a string $x \in \Sigma^{*}, \sigma(x)=m$ iff $P$ is the suffix of $x$. We can now define the transition function $\delta$ for our automaton:

$$
\delta(a, i)=\sigma\left(p_{1} \ldots p_{i} a\right)
$$

Computing the transition function. The following algorithm computes $\delta$ directly according to the definition of the $\sigma$ function.

```
Algorithm TransitionFunction \((P, \Sigma, m)\)
begin
    for \(\mathrm{q}=0\) to \(m\) do // iterate over states of the DFA
        foreach \(a \in \Sigma\) do // iterate over input characters
            \(k \leftarrow \min (m+1, q+2) ; \quad / /\) determine max length of prefix of P
            repeat // which prefix of \(P\) is a suffix of \(P[1, q] a\) ?
            \(k \leftarrow k-1 ; \quad / /\) decrease the length of prefix of P
            until \(P[1, k]\) is a suffix of \(P[1, q] a\);
            \(\delta[q, a] \leftarrow k ;\)
        end for
    end for
    return \(\delta\);
end
```

String matching. Given a computed transition function $\delta$ of the DFA $M_{P}$ constructed for $P=p_{1} \ldots p_{m}$, and an input string $S$, we can check if $S$ contains $P$ (and report all appropriate shifts) using the following straightforward algorithm:

```
Algorithm DFAMatch \((S, n \delta, m)\)
begin
    \(q \leftarrow 0 ; \quad / /\) Start in initial state
        for \(i=1\) to \(n\) do // for each input character
            \(q \leftarrow \delta(q, S[i]) ; \quad / /\) transition to a new state
            if \(q=m\) then // accepting state reached?
                output \(i-m\); // start position of \(P\) in \(S\) is \(m\) positions to the right of \(i\)
            end if
    end for
end
```

Analysis. Algorithm TransitionFunction runs in $O\left(m^{3} \cdot K\right)$ time. The outer loops repeat $m \cdot K$ times, the repeat loop runs for at most $m+1$ times, and suffix check is an $O(m)$ operation.

Algorithm DFAMatch runs in $\Theta(n)$ time, so therefore, the total running time of a single DFA-based match is $O\left(n+m^{3} \cdot K\right)$.

There is a faster way to compute $\delta$ which runs in $O(K m)$ time, which makes the overall complexity of DFA matching $O(n+K m)$. For small/fixed $K$ (for example, for our nucleotide alphabet), this reduces to $O(n+m)$.

## Knuth-Morris-Pratt Algorithm

Outline. The DFA string matching requires a precomputation of a full transition function. However, this function essentially contains "too much information". We can precompute in time $O(m)$ the essential information needed, but keep the actual matching process to $\Theta(n)$.

Example. Let $P=$ ATAT. Consider a string $S=$ ATTAATAT. Consider the initial attempt to align $P$ against the first four characters of $S$ (we show mismatches in lower case):

```
S: A T T A A T A T
P: A T a t
```

At this point, we know that $s_{2}=T$ and $s_{3}=T$. But $p_{1}=A$, so there is no need to check if $S[2,5]=P$, as we know $s_{2} \neq p_{1}$. Similarly, there is no need to check $S[3,6]=P$, because $s_{3} \neq p_{1}$. The key is that we've already acquired enough information to draw the conclusion that the next check should start at position 4:

```
S: A T T A A T A T
P: A T a t
P: A t a t
```

We learned that $s_{5}=A=p_{1}$, so next check has to come against $S[5,8]$ :

```
S: AT T A A TAT
P: A T a t
P: A t a t
P: A T AT
```

The Knuth-Morris-Pratt algorithm[3] finds a way to formalize these shortcuts and savings.

Prefix function. Given a string $P=p_{1} \ldots p_{m}$, a prefix function $\pi:\{1, \ldots, m\} \longrightarrow$ $\{0, \ldots, m-1\}$ is defined as follows:

$$
\pi[q]=\max (k \mid k<q \text { and } P[1, k] \text { is a suffix of } P[1, q])
$$

Essentially, we shift $P$ against itself and determine, when its prefixes appear as suffixes in the shift.

Example. Let $P=A A T A A T$. We construct $\pi$ as follows. For $q=1, P[1,1]=$ $A$, and the only prefix of $P$ of length less than 1 is $\epsilon$, so $\pi[1]=0$.

For $q=2, P[1,2]=A A$, the largest prefix of $P$ of length of 1 or less, which is a suffix of $P[1,2]$ is $P[1,1]=A$ :

```
P[1,2]: A A
P : A | A T A A T
```

Therefore, $\pi[2]=1$.
We continue shifting $P$ against itself (vertical bar indicates the largest substring to try):

```
P[1,3]: A A T
P : . A A | TAAT
P[1,4]: A A T A
P : A A T | A A T
P[1,5]: A A T A A
P : A A T A | A T
P[1,6]: A A T A A T
P : AATAA| T
```

This implies the following prefix function:

| $\mathrm{q}:$ | 1 | 2 | 3 | 4 | 5 | 6 |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| $\pi[\mathrm{q}]$ | 0 | 1 | 0 | 1 | 2 | 3 |
| $P[1, \pi[q]]$ | $\epsilon$ | A | $\epsilon$ | A | AA | AAT |

Computing prefix function. The following algorithm computes the prefix function:

```
Algorithm PrefixFunction \((P, m)\)
begin
    declare int \(\pi[1, . ., m]\);
    \(\pi[1] \leftarrow 0 ;\)
    \(k \leftarrow 0 ;\)
    for \(q=2\) to \(m\) do
        while \(k>0\) and \(P[k+1] \neq P[q]\) do
            \(k \leftarrow \pi[k] ;\)
        end while
        if \(P[k+1]=P[q]\) then
            \(k \leftarrow k+1 ;\)
        end if
        \(\pi[q] \leftarrow k ;\)
    end for
    return \(\pi\);
end
```

Knuth-Morris-Pratt Algorithm. The prefix function allows for the string matching algorithm to "know" how far it is possible to slide our pattern string $P$ along the input string $S$.

```
Algorithm KMP-Match \((S, P, n, m)\)
begin
    \(\pi[1, . ., m] \leftarrow \operatorname{PrefixFunction}(P, m)\);
    \(q \leftarrow 0 ; \quad\) //number of characters matched
    for \(i=1\) to \(n\) do
        while \(q>0\) and \(P[q] \neq S[i]\) do
            \(q \leftarrow \pi[q]\);
        end while
        if \(P[q+1]=S[i]\) then
            \(q \leftarrow q+1 ;\)
        end if
        if \(q=m\) then
            output \(m-i\);
        \(q \leftarrow \pi[q] ;\)
        end if
    end for
    return \(\pi\);
end
```

Analysis. KMP-Match runs in $O(n)$ plus the time it takes to construct $\pi$. PrefixFunction runs in $O(m)$, and the total running time of KMP-Match is $O(n+$ $m)$.

## Boyer-Moore Algorithm

Idea. The Boyer-Moore algorithm[1] is well designed for the following types of problems:

- Large pattern string $P$.
- Large alphabet $\Sigma$.

While the first situation is not uncommon in bioinformatics, the traditional alphabet of nucleotides $\Sigma=\{A, T, C, G\}$ is small, and Boyer-Moore algorithm might not provide the best benefit. The amino acid alphabet has 20 characters in it, for this alphabet, Boyer-Moore is a good choice.

The algorithm is based on three key ideas.

Idea 1: Right-to-left scan. All previous algorithms compared $P$ to a substring of $S$ in a left-to-right fashion, first comparing $p_{1}$ to $s_{i}$, then $p_{2}$ to $s_{i+1}$ and so on.

Boyer and Moore observed, that sometimes, one can get more information from a right-to-left comparison. E.g., consider a pattern string $P=$ AATAT and a string $S=$ AATAGTGATCG. When $P$ aligned against $S[1,5]$, we get:

```
S: A A T A G T G A T C G
P: A A T A t
```

If we attempt a left-to-right match, we will successfully match $P[1,4]$ and $S[1,4]$ before discovering $p_{5} \neq s_{5}$. We might need to consider matching $P$ against $S[2,6]$ and against $S[4,8]$.

A right-to-left match with some extra knowledge about $P$ can allow us to shift $P$ much further. In fact, $s_{5}=G$ and $G$ is not found in $P$. Therefore, no string that includes $G$ can match $P$. Thus, we can avoid comparing $P$ to $S[2,6], S[3,7], S[4,8]$ and $S[5,9]$. Our next comparison will be:

```
S: A A T A G T G A T C G
P: A A TAt
```

By comparing $s_{10}$ to $p_{5}$, we immediately reject $S[6,10]=P$. More importantly, $s_{10}=C$ and $C$ is also not in $P$, so no substring including $s_{10}$ can match $P$.

We have just discovered that $S$ does not contain $P$ using only two comparisons: $p_{5}$ vs. $s_{5}$ and $p_{5}$ vs. $s_{10}$, and scanning $P$ to determine which characters are present in it.

Idea 2: Bad character rule. Example above illustrates how Boyer-Moore algorithm proceeds when $p_{m} \neq s_{i+m}$ when comparing $P$ to $S[i, i+m]$. In a more general case, $s_{i+m}$ may be present in $P$ at one or more positons. We compute the requisite shift using a special helper function $R$.

Let $R: \Sigma \rightarrow\{1, \ldots m\}$ be defined as follows: $R(x)$ is the position of the rightmost occurrence of $x$ in $P . R(x)=0$ if $x$ does not occur in $P$.

Example. Let $\Sigma=\{A, T, C, G\}$ and let $P=$ aAtattgat. Then, $R$ is defined as follows:

$$
\begin{array}{ll}
R(A)=8 ; & R(T)=9 \\
R(C)=0 ; & R(G)=7
\end{array}
$$

$R(x)$ can be precomputed in linear time $(O(m+K))$ using the following straightforward algorithm:

```
Algorithm ComputeR \((P, m, \Sigma, \mathrm{~K})\)
begin
    for \(i=1\) to \(K\) do
        \(R\left[a_{i}\right]=0 ;\)
    end for
    for \(i=1\) to \(m\) do
        \(R[P[i]]=i\);
    end for
    return \(R\);
end
```

$R$ values are used in the bad character shift rule.
Let $P$ be aligned against $S[i, i+m]$ and let, for some $j<m, p_{j+1} \ldots p_{m}=$ $s_{i+j+1} \ldots s_{i+m}$, and $p_{j} \neq s_{i+j}$. Then, $P$ can be shifted by $\max \left(1, j-R\left(s_{i+j}\right)\right.$ places to the right, i.e., we can shift $P$ in a way, that would align $s_{i+j}$ with the rightmost appearance of the same character in $P$.

Example. Consider the situation below.

```
S: ATGACTGCATAATG
P: A GCtCA
```

$p_{4}=T$ is mismatched against $s_{7}=G$, while $p_{5}=s_{8}$ and $p_{6}=s_{9} . \quad R(T)=2$, $j=4$, and $\max (1, j-R(T))=2$, i.e, the next shift that needs to be considered, is the one, that aligns $p_{2}=G$ with $s_{7}$ :

```
S: ATGACTGCATAATG
P: AGCtCA
P: A G C T C A
```

On the other hand, if the rightmost occurence of the "bad" character overshoots the current position, we shift only by one position, as illustrated below:

```
S: A T T G A T T A A T A G
P: T cT TA
```

Here, $p_{2} \neq s_{5}$, but $p_{3}=s_{6}, p_{4}=s_{7}$ and $p_{5}=s_{8} . s_{5}=A$, but in $P, R(A)=5$, i.e., $A$ appears as the last character in $P$, and it is currently to the right of $s_{5}$. We would like to align $s_{5}$ with the closest $A$ in $P$ (if any), but we do not know where any other As are. Because of this, we prefer a conservative shift by one position to the right:

```
: A T T G A T T A A T A G
P: T c T T A
P: TCTTA
```

Idea 3. Good suffix rule. Let $P$ be matched against $S[i, i+m]$, let, for some $j<m, p_{j+1} \ldots p_{m}=s_{i+j+1} \ldots s_{i+m}$, but $p_{j} \neq s_{i+j}$.

The substring $P^{\prime}=p_{j+1} \ldots p_{m}$ is called the good suffix, because we were able to successfully match it against a substring in $S$.

If there is another occurrence of the sequence of $P^{\prime}$ in $P$, to the left of the suffix occurrence of $P^{\prime}$, we can shift $P$ to align that occurrence of $P^{\prime}$ with $s_{i+j+1} \ldots s_{i+m}$.

Example. Consider $P=$ atacgiactiac and consider the following alignment:

```
S: TGTAGATAAGTACTTACTACGTCGA
P: GTACTTACtTAC
    # * * * ^
```

$P$ is matched against $S[2,13]$ and, the good suffix of $P$ is $p_{10} p_{11} p_{12}=T A C$. $p_{9}=T$ is mismatched with $s_{1} 0=G$. We notice that $T A C$ occurs in $P$ in a substring $p_{6} p_{7} p_{8}$. We can now shift $P$ in a way that preserves the good suffix alignment against $S$ :

```
S: TGTAGATAAGTACTTACTACGTCGA
P: GTACTTACTTAC
P: G TACTTACT TAC
                # * * *
```

A strong good suffix rule can go even further, and search for an occurrence of the good suffix, prefaced with a different character. In the example above $p_{5}=p_{9}=T$ is the character that prefaces the rightmost occurence of the good suffix, and the next occurrence to the left. But we already know that $T$
positioned against $s_{1} 0$ is a mismatch, so there is no need to attempt aligning $p_{6} p_{7} p_{8}$ against $s_{11} s_{12} s_{13}$. Instead, we can look further left to see if another occurrent of $T A C$ exists, prefaced with anything other than a $T$. In our example, $p_{2} p_{3} p_{4}=T A C$ and $p_{1}=G$, so we can, indeed, attempt to align $p_{2} p_{3} p_{4}$ against $s_{11} s_{12} s_{13}$.

```
S: TGTAGATAAGTACTTACTTACGTCGA
P: GTACTTACTTAC
P: G T AC T T AC T TAC
# * * *
```

The good suffix rule heuristic can be precomputed using the following algorithm.

```
Algorithm ComputeGoodSuffixShift \((P, m)\)
begin
    \(\pi \leftarrow\) PrefixFunction \((P, m)\);
    \(P^{\prime}=\operatorname{reverse}(P)\);
    \(\pi^{\prime} \leftarrow \operatorname{PrefixFunction}\left(P^{\prime}, m\right)\);
    for \(j=0\) to \(m\) do
        \(G S[j] \leftarrow m-\pi[m] ;\)
    end for
    for \(l=1\) to \(m\) do
        \(j \leftarrow m-\pi^{\prime}[l]\);
        if \(G S[j]>l-\pi^{\prime}[l]\) then
            \(G S[j] \leftarrow l-\pi^{\prime}[l]\)
        end if
    end for
    return \(G S\);
end
```

Boyer-Moore Algorithm. Boyer-Moore Algorithm on each step tries to align $P$ and $S$ left-to-right. When a mismatch occurs, bad character rule and good suffix rule provide suggestions for the next shift. The higher of the two is chosen.

```
Algorithm Boyer-Moore \((S, n, P, m, \Sigma, K)\)
begin
    \(R \leftarrow\) ComputeR \((P, m)\);
    \(G S \leftarrow\) ComputeGoodSuffixShift \((P, m)\);
    \(s \leftarrow 0\);
    while \(s \leq n-m\) do
        \(j \leftarrow m ; \quad / /\) right-to-left scan
        while \(j>0\) and \(P[j]=S[s+j]\) do \(j \leftarrow j-1\);
        if \(j=0\) then
            output \(s\);
            \(s \leftarrow s+G S[0] ;\)
        else
            \(s \leftarrow \max (G S[j], j-R[S[s+j]]) ;\)
    end while
end
```


## Gusfield's Z Algorithm

Dan Gusfield proposed a very elegant exact string matching algorithm [2].

Z function. Given a string $S=s_{1} \ldots s_{i} \ldots s_{n}, Z_{i}(S)=k$ is the length of the longest prefix $s_{1} \ldots s_{k}$ that matches a a substring $s_{i} \ldots s_{i+k}$.

Gusfield has shown that $Z_{i}$ scores can be computed in $O(n)$ time.
Exact string matching using $Z_{i}$ scores. Let $P=p_{1} \ldots p_{m}$ be a pattern string and $S=s_{1} \ldots s_{n}$ be the query string. Consider a string $Q=P \$ S$ where $" \$ "$ is a symbol not found in either $P$ or $S$.

We find all valid shifts of $P$ in $S$ by computing $Z_{i}(Q)$ for all $i=1, \ldots n+m+1$.
We note the following:

- Because " $\$$ " is not found neither $P$ nor $S, Z_{i}(Q) \leq m$ for $i>m+1$, i.e, no substring in $S$ can have a prefix from $Q$ that goes beyond $P$.
- Any time $Z_{i}(Q)=m$ for $i>m+1$, we are documenting an occurrence of $P$ in $S$ starting with position $i-m-1$.

The algorithm then, computes $Z_{i}(Q)$ scores for $i=1 \ldots m+n+1$ and reports all positions $i$ (as $i-m-1$ ) such that $Z_{i}(Q)=m$.

## References

[1] Robert S. Boyer, J. Strother Moore (1977), A Fast String Searching Algorithm, Communications of the ACM, Vol. 20, No. 10, pp. 762-772, October 1977.
[2] Dan Gusfield (1997), Algorithms on Strings, Trees, and Sequences - Computer Science and Computational Biology, Cambridge University Press, 1997.
[3] Donald E. Knuth, James H. Morris, Vaughan R. Pratt, (1977), Fast Pattern Matching in Strings, SIAM Jorunal on Computing, Volume 6, pp. 323-350, 1977.
[4] Gonzalo Navarro, (2001), A Guided Tour to Approximate String Matching, ACM Computing Surveys, Vol. 33, No. 1, pp. 31-88, March, 2001.

