

Baeza-Yates–Perleberg Filtering Algorithm

A **filtering algorithm** for approximate string matching searches the text for factors having some **property** that satisfies the following conditions:

1. Every approximate occurrence of the pattern has this property.
2. Strings having this property are reasonably rare.
3. Text factors having this property can be found quickly.

Each text factor with the property is a **potential occurrence**, which is then **verified** for whether it is an actual approximate occurrence.

Filtering algorithms can achieve linear or even sublinear **average case time complexity**.

145

The algorithm has two phases:

Filtration: Search the text T for exact occurrences of the pattern factors P_i . Using the Aho–Corasick algorithm this takes $\mathcal{O}(n)$ time for a constant alphabet.

Verification: An area of length $\mathcal{O}(m)$ surrounding each potential occurrence found in the filtration phase is searched using the standard dynamic programming algorithm in $\mathcal{O}(m^2)$ time.

The worst case time complexity is $\mathcal{O}(m^2n)$, which can be reduced to $\mathcal{O}(mn)$ by combining any overlapping areas to be searched.

147

Many variations of the algorithm have been suggested:

- The filtration can be done with a **different multiple exact string matching algorithm**:
 - The first algorithm of this type by Wu and Manber used an extension of the Shift-And algorithm.
 - An extension of BDM achieves $\mathcal{O}(nk(\log_\sigma m)/m)$ average case search time. This is sublinear for small enough k .
 - An extension of the Horspool algorithm is very fast in practice for small k and large σ .
- Using a technique called **hierarchical verification**, the average verification time for a single potential occurrence can be reduced to $\mathcal{O}((m/k)^2)$.

A filtering algorithm by Chang and Marr has average case time complexity $\mathcal{O}(n(k + \log_\sigma m)/m)$, which is optimal.

149

4. Suffix Trees and Arrays

Let $T = T[0..n]$ be the text. For $i \in [0..n]$, let T_i denote the **suffix** $T[i..n]$. Furthermore, for any subset $C \in [0..n]$, we write $T_C = \{T_i \mid i \in C\}$. In particular, $T_{[0..n]}$ is the **set of all suffixes** of T .

Suffix tree and suffix array are **search data structures** for the set $T_{[0..n]}$.

- Suffix tree is a **compact trie** for $T_{[0..n]}$.
- Suffix array is a **ordered array** for $T_{[0..n]}$.

They support fast **exact string matching** on T :

- A pattern P has an occurrence starting at position i if and only if P is a **prefix** of T_i .
- Thus we can find all occurrences of P by a **prefix search** in $T_{[0..n]}$.

A data structure supporting fast string matching is called a **text index**.

There are numerous other applications too, as we will see later.

151

The following lemma shows the property used by the **Baeza-Yates–Perleberg algorithm** and proves that it satisfies the first condition.

Lemma 3.23: Let $P_1P_2 \dots P_{k+1} = P$ be a partitioning of the pattern P into $k + 1$ nonempty factors. Any string S with $ed(P, S) \leq k$ contains P_i as a factor for some $i \in [1..k + 1]$.

Proof. Each single symbol edit operation can change at most one of the pattern factors P_i . Thus any set of at most k edit operations leaves at least one of the factors untouched. \square

146

Let us analyze the **average case time complexity** of the verification phase.

- The best pattern partitioning is as even as possible. Then each pattern factor has length at least $r = \lfloor m/(k + 1) \rfloor$.
- The expected number of exact occurrences of a random string of length r in a random text of length n is at most n/σ^r .
- The expected total verification time is at most

$$\mathcal{O}\left(\frac{m^2(k + 1)n}{\sigma^r}\right) \leq \mathcal{O}\left(\frac{m^3n}{\sigma^r}\right).$$

This is $\mathcal{O}(n)$ if $r \geq 3 \log_\sigma m$.

- The condition $r \geq 3 \log_\sigma m$ is satisfied when $(k + 1) \leq m/(3 \log_\sigma m + 1)$.

Theorem 3.24: The average case time complexity of the Baeza-Yates–Perleberg algorithm is $\mathcal{O}(n)$ when $k \leq m/(3 \log_\sigma m + 1) - 1$.

148

Summary: Approximate String Matching

We have seen two main types of algorithms for approximate string matching:

- Basic dynamic programming time complexity is $\mathcal{O}(mn)$. The time complexity can be improved to $\mathcal{O}(kn)$ using diagonal monotonicity, and to $\mathcal{O}(n\lceil m/w \rceil)$ using bitparallelism.
- Filtering algorithms can improve average case time complexity and are the fastest in practice when k is not too large.

Other algorithms worth mentioning are those based on automata:

- Algorithms based on bit-parallel simulation of non-deterministic automata were mentioned briefly.
- A deterministic automaton can find occurrences in $\mathcal{O}(n)$ time, but the size of the automaton can be exponential in m and too big to be practical.

Similar techniques can be useful for other variants of edit distance but not always straightforwardly.

150

The set $T_{[0..n]}$ contains $|T_{[0..n]}| = n + 1$ strings of total length $\|T_{[0..n]}\| = \Theta(n^2)$. It is also possible that $\Sigma LCP(T_{[0..n]}) = \Theta(n^2)$, for example, when $T = a^n$ or $T = XX$ for any string X .

- A basic trie has $\Theta(n^2)$ nodes for most texts, which is too much. Even a leaf path compacted trie can have $\Theta(n^2)$ nodes, for example when $T = XX$ for a random string X .
- A compact trie with $\mathcal{O}(n)$ nodes and an ordered array with $n + 1$ entries have linear size.
- A compact ternary trie and a string binary search tree have $\mathcal{O}(n)$ nodes too. However, the construction algorithms and some other algorithms we will see are not straightforward to adapt for these data structures.

Even for a compact trie or an ordered array, we need a **specialized construction algorithm**, because any general construction algorithm would need $\Omega(\Sigma LCP(T_{[0..n]}))$ time.

152


Suffix links are the same as Aho–Corasick [failure links](#) but Lemma 4.4 ensures that $\text{depth}(\text{slink}(u)) = \text{depth}(u) - 1$. This is not the case for an arbitrary trie or a compact trie.

Suffix links are stored for compact trie nodes only, but we can define and compute them for any locus (u, d) :

```

slink(u, d)
(1) v ← slink(parent(u))
(2) while depth(v) < d - 1 do
(3)   v ← child(v, T[start(u) + depth(v) + 1])
(4) return (v, d - 1)

```



161

The creation of a new node on line (6) never happens in a fully constructed suffix tree, but during the brute force algorithm the necessary node may not exist yet:

- If a new internal node u_i was created during the insertion of the suffix T_i , there exists an earlier suffix T_j , $j < i$ that branches at u_i into a different direction than T_i .
- Then $\text{slink}(u_i)$ represents a prefix of T_{j+1} and thus exists at least as a locus on the path labelled T_{j+1} . However, it may be that it does not become a branching node until the insertion of T_{i+1} .
- In such a case, $\text{ComputeSlink}(u_i)$ creates $\text{slink}(u_i)$ a moment before it would otherwise be created by the brute force construction.

163

Algorithm 4.5: McCreight

Input: text $T[0..n]$ ($T[n] = \$$)

Output: suffix tree of T : *root*, *child*, *parent*, *depth*, *start*, *slink*

```

(1) create new node root; depth(root) ← 0; slink(root) ← root
(2) u ← root; d ← 0 // (u, d) is the active locus
(3) for i ← 0 to n do // insert suffix T_i
(4)   while d = depth(u) and child(u, T[i + d]) ≠ ⊥ do
(5)     u ← child(u, T[i + d]); d ← d + 1
(6)   while d < depth(u) and T[start(u) + d] = T[i + d] do d ← d + 1
(7)   if d < depth(u) then // (u, d) is in the middle of an edge
(8)     u ← CreateNode(u, d)
(9)   CreateLeaf(i, u)
(10)  if slink(u) = ⊥ then ComputeSlink(u)
(11)  u ← slink(u); d ← d - 1

```

165

There are other linear time algorithms for suffix tree construction:

- Weiner's algorithm was the first. It inserts the suffixes into the tree in the opposite order: T_n, T_{n-1}, \dots, T_0 .
- Ukkonen's algorithm constructs suffix tree first for $T[0..1]$ then for $T[0..2]$, etc.. The algorithm is structured differently, but performs essentially the same tree traversal as McCreight's algorithm.
- All of the above are linear time only for constant alphabet size. Farach's algorithm achieves linear time for an integer alphabet of polynomial size. The algorithm is complicated and unpractical.

167

The same idea can be used for computing the suffix links during or after the brute force construction.

```

ComputeSlink(u)
(1) d ← depth(u)
(2) v ← slink(parent(u))
(3) while depth(v) < d - 1 do
(4)   v ← child(v, T[start(u) + depth(v) + 1])
(5) if depth(v) > d - 1 then // no node at (v, d - 1)
(6)   v ← CreateNode(v, d - 1)
(7) slink(u) ← v

```

The procedure $\text{CreateNode}(v, d - 1)$ sets $\text{slink}(v) = \perp$.


The algorithm uses the suffix link of the parent, which must have been computed before. Otherwise the order of computation does not matter.

162

McCreight's Algorithm

McCreight's suffix tree construction is a simple modification of the brute force algorithm that computes the suffix links during the construction and uses them as [short cuts](#):

- Consider the situation, where we have just added a leaf w_i representing the suffix T_i as a child to a node u_i . The next step is to add w_{i+1} as a child to a node u_{i+1} .
- The brute force algorithm finds u_{i+1} by traversing from the root. McCreight's algorithm takes a short cut to $\text{slink}(u_i)$.



- This is safe because $\text{slink}(u_i)$ represents a prefix of T_{i+1} .

164

Theorem 4.6: Let T be a string of length n over an alphabet of constant size. McCreight's algorithm computes the suffix tree of T in $\mathcal{O}(n)$ time.

Proof. Insertion of a suffix T_i takes constant time except in two points:

- The while loops on lines (4)–(6) traverse from the node $\text{slink}(u_i)$ to u_{i+1} . Every round in these loops increments d . The only place where d decreases is on line (11) and even then by one. Since d can never exceed n , the total time on lines (4)–(6) is $\mathcal{O}(n)$.
- The while loop on lines (3)–(4) during a call to $\text{ComputeSlink}(u_i)$ traverses from the node $\text{slink}(\text{parent}(u_i))$ to $\text{slink}(u_i)$. Let d'_i be the depth of $\text{parent}(u_i)$. Clearly, $d'_{i+1} \geq d'_i - 1$, and every round in the while loop during $\text{ComputeSlink}(u_i)$ increases d'_{i+1} . Since d'_i can never be larger than n , the total time in the loop on lines (3)–(4) in ComputeSlink is $\mathcal{O}(n)$. □

166