

case (two copies of a long string). See references mentioned above for worst case solutions.

```
void standard_lookup(char* s, int* suf, int N){
    int* i = find_first_suf(s, suf, N);
    int* j = find_last_suf(s, suf, N);
    for (int* k=i; k<=j; k++) output(*k);}

int* find_first_suf(char* s, int* suf, int N) {
    int len = strlen(s);
    int* high = suf + N;
    while (suf + 2 < high) {
        int* mid = suf + (high-suf)/2;
        int c = strncmp(s, corpus + *mid, len);
        if (c == 0) high = mid+1;
        else if (c < 0) high = mid;
        else suf = mid;}
    for ( ; suf < high; suf++)
        if (strncmp(s, corpus + *suf, len) == 0)
            return suf;
    return NULL;} // not found
```

2 K-Best Suffix Arrays

K-best suffix arrays are like standard suffix arrays, except there are two orders instead of one. In addition to lexicographic order, we assume a figure of merit, which we will refer to as popularity. For example, the popularity of a string could be its frequency in a search log. The code below assumes that the corpus is a sequence of strings that comes pre-sorted by popularity, and then the popularities have been stripped off. These assumptions make it very easy to compare two strings by popularity. All *popcomp* has to do is to compare the two positions in the corpus.¹

The *make_kbest_suf* program below is similar to the *make_standard_suf* program above except we now sort by the two orders at alternating depths in the tree. First we sort lexicographically and then we sort by popularity and so on, using a construction similar to KD-Trees (Bentley, 1975). The code below is simple to describe (though there are more efficient implementations that avoid unnecessary qsorts).

```
int* make_kbest_suf() {
    int N = strlen(corpus);
    int* suf = (int*)malloc(N * sizeof(int));
```

¹ With a little extra book keeping, one can keep a table on the side that makes it possible to map back and forth between popularity rank and the actual popularity. This turns out to be useful for some applications.

```
for (int i=0; i<N; i++) suf[i]=i;
process(suf, suf+N, 0);
return suf;}
```

```
void process(int* start, int* end, int depth) {
    int* mid = start + (end - start)/2;
    if (end <= start+1) return;
    qsort(start, end-start, sizeof(int),
        (depth & 1) ? popcomp : lexcomp);
    process(start, mid, depth+1);
    process(mid+1, end, depth+1);}
```

```
int popcomp(int* a, int* b) {
    if (*a > *b) return 1;
    if (*a < *b) return -1;
    return 0;}
```

2.1 K-Best Suffix Array Lookup

To find the k-best matches for a particular substring *s*, we do what we would normally do for standard suffix arrays on lexicographic splits. However, on popularity splits, we search the more popular half first and then we search the less popular half, if necessary.

An implementation of *kbest_lookup* is given below. *D* denotes the depth of the search thus far. *Kbest_lookup* is initially called with *D* of 0. *Propose* maintains a heap of the k-best matches found thus far. *Done* returns true if its argument is less popular than the *k*th best match found thus far.

```
void kbest_lookup(char* s, int* suf, int N, int D){
    int* mid = suf + N/2;
    int len = strlen(s);

    if (N==1 && strncmp(s, corpus+*suf, len)==0)
        propose(*suf);
    if (N <= 1) return;

    if (D&1) { // popularity split
        kbest_lookup(s, suf, mid-suf, D+1);
        if (done(*mid)) return;
        if (strncmp(s, corpus + *mid, len) == 0)
            propose(*mid);
        kbest_lookup(s, mid+1, (suf+N)-mid-1,
            D+1);}

    else { // lexicographic split
        int c = strncmp(s, corpus + *mid, len);
        int n = (suf+N)-mid-1;
        if (c < 0) kbest_lookup(s, suf, mid-suf, D+1);
        else if (c > 0) kbest_lookup(s, mid+1, n, D+1);
        else { kbest_lookup(s, suf, mid-suf, depth+1);
            propose(*mid);
            kbest_lookup(s, mid+1, n, D+1); }}}
```

2.2 A Short Example: To be or not to be

Suppose we were given the text, “to be or not to be.” We could then generate the following dictionary with frequencies (popularities).

Popularity	Word
2	to
2	be
1	or
1	not

The dictionary is sorted by popularity. We treat the second column as an $N=13$ byte corpus (with underscores at record boundaries): to_be_or_not_

Standard		K-Best	
suf	corpus + suf[i]	suf	corpus + suf[i]
12	_	2	_be_or_not_
2	_be_or_not_	3	be_or_not_
8	_not_	4	e_or_not_
5	_or_not_	5	_or_not_
3	be_or_not_	8	_not_
4	e_or_not_	12	_
9	not_	9	not_
1	o_be_or_not_	1	o_be_or_not_
6	or_not_	6	or_not_
10	ot_	0	to_be_or_not_
7	r_not_	7	r_not_
11	t_	10	ot_
0	to_be_or_not_	11	t_

The standard suffix array is the 1st column of the table above. For illustrative convenience, we show the corresponding strings in the 2nd column. Note that the 2nd column is sorted lexicographically.

The k-best suffix array is the 3rd column with the corresponding strings in the 4th column. The first split is a lexicographic split at 9 (“not_”). On both sides of that split we have a popularity split at 5 (“_or_not_”) and 7 (“r_not_”). (Recall that relative popularity depends on corpus position.) Following there are 4 lexicographic splits, and so on.

If k-best lookup were given the query string $s = \text{“o,”}$ then it would find 1 (o_be_or_not_), 6 (or_not_) and 10 (ot_) as the best choices (in that order). The first split is a lexicographic split. All

the matches are below 9 (not_). The next split is on popularity. The matches above this split (1&6) are as popular as the matches below this split (10).

It is often desirable to output matching records (rather than suffixes). Records are output in popularity order. The actual popularity can be output, using the side table mentioned in footnote 1:

Popularity	Record
2	to
1	or
1	not

2.3 Time and Space Complexity

The space requirements are the same for both standard and k-best suffix arrays. Both indexes are permutations of the same suffixes.

The time requirements are quite different. Standard suffix arrays were designed to find all matches, not the k-best. Standard suffix arrays can find all matches in $O(\log N)$ time. However, if we attempt to use standard suffix arrays to find the k-best, something they were not designed to do, then it could take a long time to sort through the worst case (an embarrassment of riches with lots of matches). When the query matches every string in the dictionary, standard suffix arrays do not help us find the best matches. K-best suffix arrays were designed to handle an embarrassment of riches, which is quite common, especially when the substring s is short. Each popularity split cuts the search space in half when there are lots of lexicographic matches.

The best case for k-best suffix arrays is when the popularity splits always work in our favor and we never have to search the less popular half. The worst case is when the popularity splits always fail, such as when the query string s is not in the corpus. In this case, we must always check both the popular half and the unpopular half at each split, since the failure to find a lexicographic match in the first tells us nothing about the existence of matches in the second.

Asymptotically, k-best lookup takes between $\log N$ and \sqrt{N} time. To see this complexity result, let $P(N)$ be the work to process N items starting with a popularity splits and let $L(N)$ be the work to process N items starting with a lexicographic splits.

Thus,

$$P(N) = \alpha L(N/2) + C_1$$

$$L(N) = P(N/2) + C_2$$

where $\alpha = 2-p$, when p is the probability that the popular half contains sufficient matches. α lies between 1 (best case) and 2 (worst case). C_1 and C_2 are constants. Thus,

$$P(N) = \alpha P(N/4) + C \quad (1)$$


where $C = C_1 + \alpha C_2$. Using the master method (Cormen *et al*, 2001), $P(N) = O(\log_2 N)$ in the best case ($\alpha=1$). In the worst case ($\alpha=2$), $P(N) = O(\sqrt{N})$. In general, for $\alpha > 1$, $P(N) = O(N^{(\log_2 \alpha)/2})$.

In practical applications, we expect popularity splits to work more often than not, and therefore we expect the typical case to be closer to the best case than the worst case.

3 Empirical Study

The plot below shows the k-best lookup time as a function of square root of corpus size. We extracted sub-corpora from a 150 MB collection of 8M queries, sorted by popularity, according to the logs from Microsoft www.live.com. All experiments were performed on a Pentium 4, 3.2GHz dual processor machine with enough memory to avoid paging.

The line of diamonds shows the worst case, where we the query string is not in the index. Note that the diamonds fit the regression line quite well, confirming the theory in the previous section: The worst case lookup is $O(\sqrt{N})$.



To simulate a more typical scenario, we constructed random samples of queries by popularity, represented by squares in the figure. Note that the squares are well below the line, demonstrating that these queries are considerably easier than the worst case.

K-best suffix arrays have been used in auto-complete applications (Church and Thiesson, 2005). The triangles with the fastest lookup times demonstrate the effectiveness of the index for this application. We started with the random sample above, but replaced each query q in the sample with a substring of q (of random size).

4 Conclusion

A new data structure, *k-best suffix arrays*, was proposed. K-best suffix arrays are sorted by two orders, lexicographic and popularity, which make it convenient to find the most popular matches, especially when there are lots of matches. In many applications, such as the web, there are often embarrassments of riches (lots of matches).

Lookup time varies from $\log N$ to \sqrt{N} , depending on the effectiveness of the popularity splits. In the best case (e.g., very short query strings that match nearly everything), the popularity splits work nearly every time and we rarely have to search the less popular side of a popularity split. In this case, the time is close to $\log N$. On the other hand, in the worst case (e.g., query strings that match nothing), the popularity splits never work, and we always have to search both sides of a popularity split. In this case, lookup time is \sqrt{N} . In many cases, popularity splits work more often than not, and therefore, performance is closer to $\log N$ than \sqrt{N} .

References

- Jon Louis Bentley. 1975. Multidimensional Binary Search Trees Used for Associative Searching, *Communications of the ACM*, 18:9, pp. 509-517.
- Kenneth Church and Bo Thiesson. 2005. The Wild Thing, *ACL*, pp. 93-96.
- Udi Manber and Gene Myers. 1990. Suffix Arrays: A New Method for On-line String Searches, *SODA*, pp. 319-327.
- Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. 2001. Introduction to Algorithms, Second Edition. *MIT Press and McGraw-Hill*, pp.73-90.

It contains well written, well thought and well explained computer science and programming articles, quizzes and practice/competitive programming/company interview Questions. We have discussed Naive algorithm for construction of suffix array. The Naive algorithm is to consider all suffixes, sort them using a $O(n \log n)$ sorting algorithm and while sorting, maintain original indexes. Time complexity of the Naive algorithm is $O(n^2 \log n)$ where n is the number of characters in the input string. In this post, a $O(n \log n)$ algorithm for suffix array construction is discussed. Let us first discuss a $O(n * \log n * \log n)$ algorithm for simplicity. The idea is to use the fact that strings that are to be sorted are suffixes of a single string. Suffix arrays are a programming contest approach. Adrian Vladu and Cosmin Negru's. Summary. To get an algorithm of good complexity, this step must be done efficiently. We may use a segment tree, walk with k from 1 to n and delete from the tree the elements smaller than k , while updating the tree so that it will answer queries like "what is the greatest distance between two consecutive elements in the structure. Suffix arrays: How to compute them fast with Python". By Louis Abraham. You can download the notebook from https://louisabraham.github.io/notebooks/suffix_arrays.ipynb. You can comment there on HN. Yes, I'm using HN as a comment system. For example, the suffix array of 'banana' is [3, 2, 5, 1, 4, 0]. We can see that the first two suffixes (ranks 0 and 1) are the ones that begin with the letter 'a'. There is an equivalent definition of suffix arrays: `sa_alternative[i]` is the index of the suffix at the i -th position in the sorted list of all suffixes. In [2]: `def suffix_array_alternative_naive(s): return [rank for suffix, rank in sorted((s[i:], i) for i in range(len(s)))]. suffix_array_alternative_naive('banana'). Out[2]: [5, 3, 1, 0, 4, 2].`