

APPROXIMATION ALGORITHMS FOR THE SHORTEST COMMON SUPERSTRING PROBLEM

Jonathan S. Turner*
Computer Science Department
Washington University, St. Louis

Abstract. The object of the shortest common superstring problem (SCS) is to find the shortest possible string that contains every string in a given set as substrings. As the problem is NP-complete, approximation algorithms are of interest. The value of an approximate solution to SCS is normally taken to be its length, and we seek algorithms that make the length as small as possible. A different measure is given by the sum of the overlaps between consecutive strings in a candidate solution. When considering this measure, the object is to find solutions that make it as large as possible. These two measures offer different ways of viewing the problem. While the two viewpoints are equivalent with respect to optimal solutions, they differ with respect to approximate solutions. We describe several approximation algorithms that produce solutions that are always within a factor of two of optimum with respect to the overlap measure. We also describe an efficient implementation of one of these, using McCreight's compact suffix tree construction algorithm. The worst-case running time is $O(m \log n)$ for small alphabets, where m is the sum of the lengths of all the strings in the set and n is the number of strings. For large alphabets, the algorithm can be implemented in $O(m \log m)$ time by using Sleator and Tarjan's lexicographic splay tree data structure.

Key words. shortest common superstring, approximation algorithms, longest path problem, traveling salesman problem, suffix trees

1. Introduction. Let $s_1 = a_1 \dots a_r$ and $s_2 = b_1 \dots b_s$ be strings over some finite alphabet Σ . We say that s_1 is a *substring* of s_2 if there is an integer $i \in [0, s - r]$ such that $a_j = b_{i+j}$ for $1 \leq j \leq r$. We also say in this case that s_2 is a *superstring* of s_1 .

An instance of the *shortest common superstring problem* (SCS) is a set of strings $S = \{s_1, \dots, s_n\}$ over a finite alphabet Σ . The object of the problem is to find a minimum length string that is a superstring of every $s_i \in S$. We let $\phi^*(S)$ denote the length of a minimum length superstring.

EXAMPLE: If $S = \{\text{egiach}, \text{bfgiak}, \text{hfdegi}, \text{iakhfd}, \text{fgiakh}\}$, the string bfgiakhfdegiach is a solution of length 15.

We say that a set of strings is *substring free* if no string in the set is a substring of any other. We will generally limit our attention to substring free sets. This involves no loss of generality, since any set of strings has a unique substring free subset which has the same solutions as the original set.

We have presented the problem in the conventional way, with the object being to minimize the solution length. It is useful to consider an alternative viewpoint as

^{0*} The research described here was supported by the National Science Foundation (grant #DCR-8409435) and the National Institutes of Health (grant #RR-01380)

well. One can view the object of the problem as being to find an ordering of the strings that maximizes the amount of overlap between consecutive strings. To make this precise we need a few definitions.

Let $s_1 = a_1 \dots a_r$ and $s_2 = b_1 \dots b_s$ be strings. We define

$$\psi(s_1, s_2) = \max \{k \geq 0 \mid a_{r-k+i} = b_i \ 1 \leq i \leq k\}$$

If $\psi(s_1, s_2) = k$ then $s_1 \circ s_2$ is defined to be the string $a_1 \dots a_r, b_{k+1} \dots b_s$. We note that if s_1, s_2, s_3 are strings, none of which is a substring of another, then $s_1 \circ (s_2 \circ s_3) = (s_1 \circ s_2) \circ s_3$; that is, the overlapping operation is associative for substring free sets. Consequently, we may write $s_1 \circ s_2 \circ \dots \circ s_n$ with no ambiguity.

Let π be a permutation on $\{1, \dots, n\}$. We will usually write π_i for $\pi(i)$. We define

$$\psi_\pi(s_1, \dots, s_n) = \sum_{i=1}^{n-1} \psi(s_{\pi_i}, s_{\pi_{i+1}})$$

and $\phi_\pi(s_1, \dots, s_n) = |s_{\pi_1} \circ \dots \circ s_{\pi_n}|$. Note that for any instance $S = (s_1, \dots, s_n)$ of SCS,

$$\phi_\pi(S) = \|S\| - \psi_\pi(S)$$

where $\|S\| = \sum_{i=1}^n |s_i|$. In particular,

$$\phi^*(S) = \|S\| - \psi^*(S) \quad \text{where} \quad \psi^*(s_1, \dots, s_n) = \max_{\pi} \psi_\pi(s_1, \dots, s_n)$$

Hence, we can view the object of the SCS problem as being to find a mapping π that maximizes ψ_π .

Let A be an algorithm for SCS which given a collection of strings $S = (s_1, \dots, s_n)$ produces a mapping $\pi = \pi_A(S)$. We define $\psi_A(S) = \psi_\pi(S)$ and $\phi_A(S) = \phi_\pi(S)$.

SCS was shown to be NP-complete by Maier and Storer in [8]. Another, and more elegant proof appears in [3] and [4]. One obvious application for the problem is data compression. Storer and Szymanski [15] for example, consider a fairly general model of data compression which includes SCS as an important special case. See also [9]. Another application is to DNA sequencing. SCS is one of the simplest models for the problem of recovering DNA sequencing information from experimental data [5, 12, 14]. To our knowledge the only approximation algorithm to be discussed in the literature is a simple greedy algorithm which is treated briefly by Gallant in [4]. Gallant claims that for this algorithm, which we refer to as SGREEDY, $\phi_{\text{SGREEDY}}(S) \leq (3/2)\phi^*(S)$ for all collections of strings S . We show that this is not in fact true by displaying a set of strings S for which $\phi_{\text{SGREEDY}}(S) \approx 2\phi^*(S)$. We have found no worse example problem than this, but have also been unsuccessful in proving an upper bound on the performance of this algorithm in terms of the length measure. On the other hand, we do show that $\psi^*(S) \leq 2\psi_{\text{SGREEDY}}(S)$.

In section 2 we relate SCS to the longest path problem (LPP) in graphs by describing a transformation from SCS to LPP that preserves solution values with respect to the overlap measure. We then construct three approximation algorithms for LPP, two based on matching and the third a greedy heuristic. By virtue of the transformation

from SCS, all three are also approximation algorithms for SCS. We show that the greedy heuristic for LPP always produces solutions within a factor of three of the optimum value. In section 3, we show that the instances of LPP that result from our transformation from SCS have a special structure that allows us to obtain a tighter bound. We also describe an efficient implementation of this greedy algorithm for strings using a compact representation of suffix trees. In section 4, we relate SCS to the traveling salesman problem (TSP) by another transformation that preserves solution values, this time with respect to the length measure. The instances of TSP arising from this transformation are asymmetric, but satisfy the triangle inequality. There are no approximation algorithms known for this problem with provably good worst-case performance, nor have we succeeded in finding any. Nevertheless, this transformation means that if such an algorithm is found, it can be used for SCS as well as TSP. If on the other hand, it turns out that approximating this version of TSP is hard, then any approximation algorithm for SCS, will have to make use of special structural properties present in the instances of TSP that arise from this transformation.

2. SCS and the Longest Path Problem. In this section we relate SCS to the *longest path problem* (LPP) in graphs. An instance of the longest path problem is a complete directed graph $G = (V, E)$ with each edge (u, v) having a non-negative integer length $\ell(u, v)$. The *length of a path* p in G is defined to be the sum of the lengths of its edges and is denoted $\lambda_p(G, \ell)$. The object of the longest path problem is to find a Hamiltonian path p (that is a path including every vertex) in G that maximizes $\lambda_p(G, \ell)$. The length of such a longest path is denoted $\lambda^*(G, \ell)$.

Let $S = (s_1, \dots, s_n)$ be an instance of SCS. We define $\text{LPP}(S)$ to be an instance (G, ℓ) of LPP with

$$\begin{aligned} V &= \{u_1, \dots, u_n\} & E &= V \times V \\ \ell(u_i, u_j) &= \psi(s_i, s_j) & & 1 \leq i, j \leq n, i \neq j \end{aligned}$$

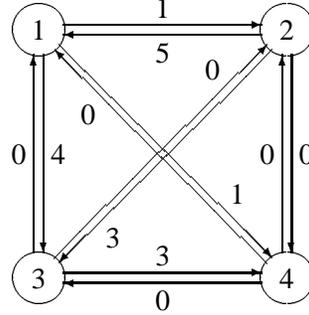
An example of this transformation is shown in Figure 1.

Let π be a permutation on $\{1, \dots, n\}$. We can view π as defining a Hamiltonian path $u_{\pi_1}, \dots, u_{\pi_n}$ in G . We let $\lambda_\pi(G, \ell)$ denote the length of this path. We now state a trivial, but useful theorem.

THEOREM 2.1. *Let $S = (s_1, \dots, s_n)$ be an instance of SCS, $(G, \ell) = \text{LPP}(S)$ and let π be any permutation on $\{1, \dots, n\}$. $\lambda_\pi(G, \ell) = \psi_\pi(S)$. In particular, $\lambda^*(G, \ell) = \psi^*(S)$.*

The theorem implies that any approximation algorithm for LPP is an approximation algorithm for SCS with respect to the overlap measure. In the remainder of this section, we present three simple approximation algorithms for LPP.

2.1. Matching Algorithm. A *matching* in a graph $G = (V, E)$ is a set of edges, no two of which share a common vertex. A maximum matching in a graph with



$$S = \{cbadef, fcbade, adefcd, fcdafeb\}$$

Figure 1: Example of transformation from SCS to LPP

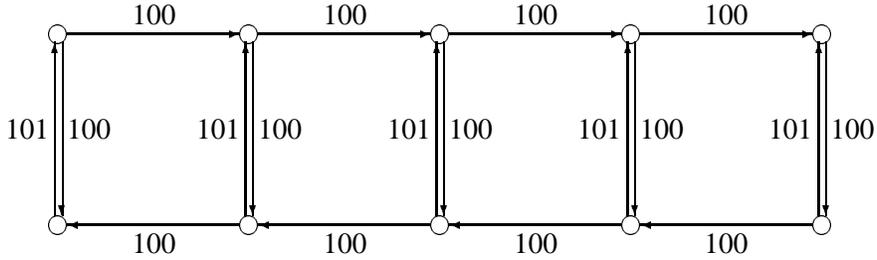


Figure 2: Worst-Case Example for Theorem 2.2

edge lengths $\ell(e)$ is a matching M that maximizes $\ell(M)$. We define $\mu^*(G, \ell) = \max_M \ell(M)$ to be the value of a maximum matching. There are algorithms for finding maximum matchings having running times of $O(n^3)$ (where $n = |V|$) [16].

Our first algorithm for LPP is based on the observation that any matching for an instance (G, ℓ) of LPP can be extended to a path (since G is assumed to be complete) and a maximum matching must have total length at least half that of a longest path. (Recall that we are restricting attention to non-negative weights.)

THEOREM 2.2. *If $(G = (V, E), \ell)$ is an instance of LPP then $\lambda^*(G, \ell) \leq 2\mu^*(G, \ell)$.*

Proof. Let P be a set of edges defining any Hamiltonian path. Let Q be obtained by taking alternate edges from P and let $R = P - Q$. Both Q and R are matchings. The sum of the lengths of the edges in Q is $\leq \mu^*(G, \ell)$. Similarly, the sum of the lengths of the edges in R is $\leq \mu^*(G, \ell)$. Hence, $\ell(P) \leq 2\mu^*(G, \ell)$ and since this holds for all paths P , it follows that $\lambda^*(G, \ell) \leq 2\mu^*(G, \ell)$. \square

Remark. There are instances (G, ℓ) of LPP for which $\lambda^*(G, \ell)$ approaches $2\mu^*(G, \ell)$. Figure 2 shows a graph for which $\lambda^*(G, \ell) = 901$ and $\mu^*(G, \ell) = 505$. (The edges not explicitly shown have length 0.) The example is easily extended to give graphs for which the ratio λ^*/μ^* is arbitrarily close to 2.

```

function edgeset MATCH(digraph  $G = (V, E)$ , edgelengths  $\ell$ )
  edgeset  $P, M$ ;
   $P := \emptyset$ ;
  do  $E \neq \emptyset \rightarrow$ 
     $M := \text{MAXMATCH}(G, \ell)$ ;
     $P := P \cup M$ ;
    for  $(u, v) \in M \rightarrow$ 
      Delete from  $G$ , all edges of the form  $(u, x)$  or  $(y, v)$ ;
      Collapse  $u$  and  $v$  into a single vertex;
    rof;
  od;
  return  $P$ ;
end;

```

Figure 3: Matching Algorithm for LPP

Theorem 2.2 provides the basis for our first approximation algorithm shown in Figure 3. The procedure MATCH starts by finding a maximum matching in G , then removes edges that are ruled out by the selected edges, collapses the selected edges into single vertices and then repeats the process on the new graph. To see that the algorithm does construct a Hamiltonian path, note the following: (1) the edge eliminations ensure that the set P never contains two edges leaving a common vertex or entering a common vertex, (2) the collapsing of edges into single vertices prevents creation of cycles and (3) since the original graph is assumed to be complete, the algorithm will halt only when a complete Hamiltonian path has been constructed. An example illustrating the operation of the algorithm is given in Figure 4.

Theorem 2.2 implies that $\psi^*(G, \ell) \leq 2\psi_{\text{MATCH}}(G, \ell)$ for any instance (G, ℓ) of LPP. This cannot be improved, as can be seen by considering the operation of MATCH on the graph in Figure 2. The running time of MATCH is determined primarily by the matching algorithm used. Assuming a matching algorithm that runs in $O(n^3)$ time, we get a running time of $O(n^3 \log n)$ for MATCH.

2.2. Directed Matching Algorithm. A *directed matching* in a digraph $G = (V, E)$ is a set of edges, no two of which enter a common vertex and no two of which leave a common vertex. In other words, it is a subgraph of G comprising a collection of disjoint paths and cycles. A maximum directed matching in a graph G with edge lengths $\ell(e)$ is a directed matching M that maximizes $\ell(M)$. We define $\delta^*(G, \ell) = \max_M \ell(M)$ to be the value of a maximum directed matching (where in this case, M ranges over all directed matchings of G). There are algorithms for finding a maximum directed matchings having running times of $O(n^{5/2})$ [16].

Given any matching M , let M^- be a subset of M obtained by discarding a least cost edge from each cycle in M . Our next algorithm for LPP is based on the

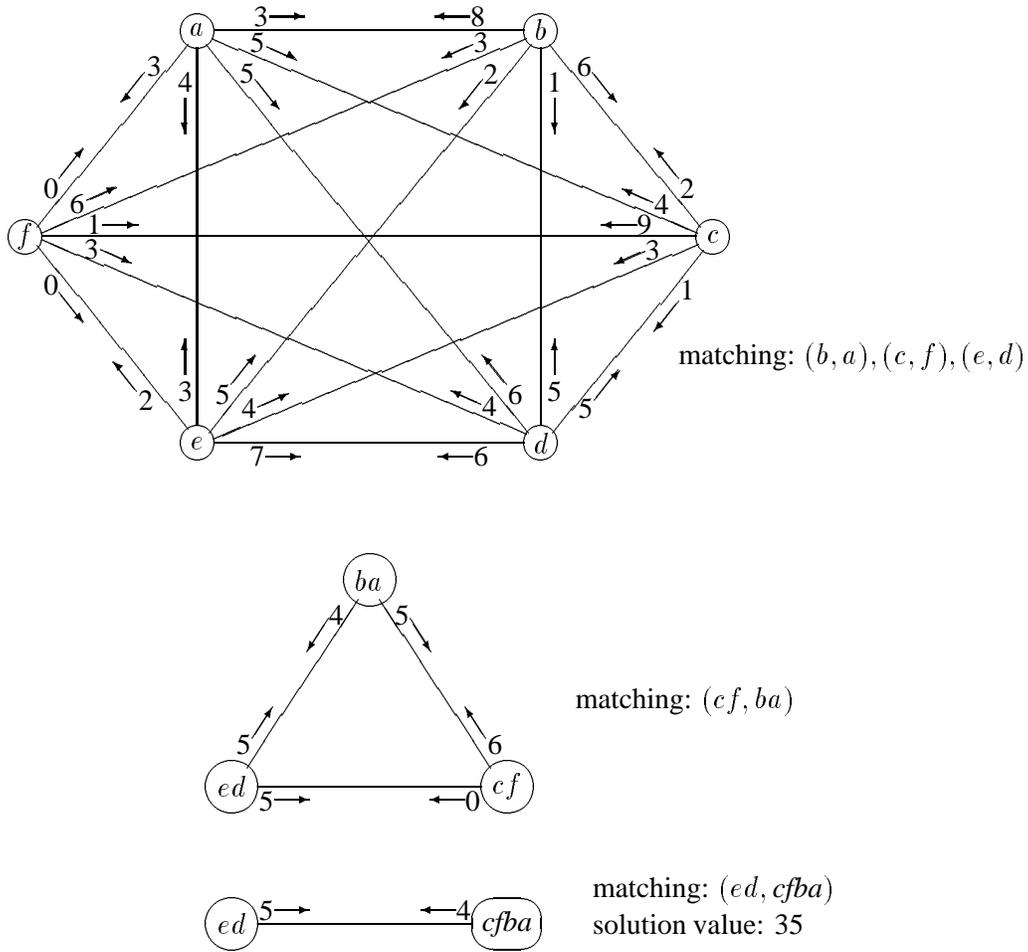


Figure 4: Example of Algorithm MATCH

observation contained in the next theorem.

THEOREM 2.3. *Let $(G = (V, E), \ell)$ be an instance of LPP, let M be a maximum directed matching of G and let k be the minimum number of edges in any cycle defined by M . $\lambda^*(G, \ell) \leq \frac{k}{k-1} \ell(M^-)$. In particular, $\lambda^*(G, \ell) \leq 2\ell(M^-)$.*

Proof. Let P be a set of edges defining a path and let M be a maximum directed matching. Notice that P is a directed matching and hence $\ell(P) \leq \ell(M)$. Let C be a cycle in M with h edges and let C^- be a path obtained by discarding a minimum length edge from C .

$$\ell(C) \leq \frac{h}{h-1} \ell(C^-) \leq \frac{k}{k-1} \ell(C^-)$$

```

function edgeset DIMATCH(digraph  $G = (V, E)$ , edgelengths  $\ell$ )
  edgeset  $P, M$ ;
   $P := \emptyset$ ;
  do  $E \neq \emptyset \rightarrow$ 
     $M := \text{MAXDIMATCH}(G, \ell)$ ;
     $M^- := M - \text{one least cost edge from each cycle of } M$ ;
     $P := P \cup M^-$ ;
    for each path  $(u_1, \dots, u_r) \in M^- \rightarrow$ 
      Delete from  $G$ , all edges of the form  $(u_i, x), 2 \leq i \leq r$ ;
      Delete from  $G$ , all edges of the form  $(x, u_i), 1 \leq i \leq r - 1$ ;
      Delete from  $G$  the edge  $(u_r, u_1)$ , if present;
      Collapse the path into a single vertex;
    rof;
  od;
  return  $P$ ;
end;

```

Figure 5: Directed Matching Algorithm for LPP

Also, for every path $R \in M$, $\ell(R) \leq \frac{k}{k-1}\ell(R)$. Summing over all paths and cycles in M yields $\ell(M) \leq \frac{k}{k-1}\ell(M^-)$. Since this is true for all paths P and since $\ell(P) \leq \ell(M)$, $\lambda^*(G, \ell) \leq \frac{k}{k-1}\ell(M^-)$. \square

Remark. There are instances (G, ℓ) of LPP for which $\lambda^*(G, \ell)$ approaches $2\ell(M^-)$. Consider for example, the graph shown in Figure 2. For this graph $\lambda^*(G, \ell) = 901$ and the optimum directed matching consists of five cycles each having two edges and length 201. When the cycles are broken, we have $\ell(M^-) = 505$. The example is easily extended to give graphs for which the ratio $\lambda^*(G, \ell)/\ell(M^-)$ is arbitrarily close to 2.

We note that $\delta^*(G, \ell) \geq \lambda^*(G, \ell)$. Hence, it provides a measure of how close a given solution is to optimal. We expect that the solutions obtained by breaking cycles will often be much closer to optimal than the bound in the theorem implies.

Theorem 2.3 provides the basis for our next approximation algorithm for LPP, shown in Figure 5. This algorithm constructs a maximum directed matching M in G , then breaks all the cycles in M and constructs a new graph in which the paths of M correspond to vertices. It then proceeds by finding a maximum directed matching in the new graph, continuing in this fashion until a Hamiltonian path in the original graph has been found. To verify that the algorithm does construct a Hamiltonian path, it suffices to note the following: (1) the edge eliminations ensure that the set P never contains two edges leaving a common vertex or entering a common vertex, (2) cycles formed are explicitly broken and the broken edges removed from the graph and (3) since the original graph is assumed to be complete, the algorithm will halt only

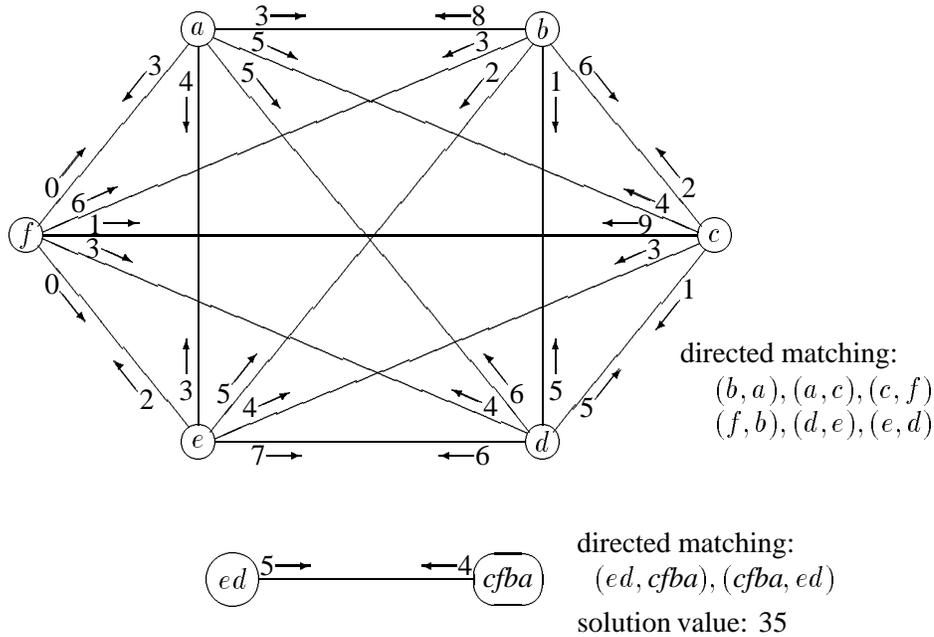


Figure 6: Example of Algorithm DIMATCH

when a complete Hamiltonian path has been constructed. An example illustrating the operation of the algorithm is given in Figure 6.

Theorem 2.3 implies that $\psi^*(G, \ell) \leq 2\psi_{\text{DIMATCH}}(G, \ell)$ for any instance (G, ℓ) of LPP. This cannot be improved, as can be seen by considering the operation of MATCH on the graph in Figure 2. The running time of DIMATCH is determined primarily by the directed matching algorithm used. Assuming an algorithm that runs in $O(n^{5/2})$ time, we get a running time of $O(n^{5/2} \log n)$ for DIMATCH.

DIMATCH is essentially an adaptation of an algorithm for the asymmetric traveling salesman problem (TSP) described by Karp in [6]. Karp's algorithm has poor worst-case performance for TSP, but performs well in a probabilistic sense for instances in which inter-city distances are selected uniformly on the interval $[0, 1]$. We have simply adapted his algorithm to the longest path problem (simplifying it slightly in the process), and observed that its worst-case performance is provably good in this context.

2.3. Greedy algorithm. The algorithms considered above are both fairly complicated and time consuming because they require the calculation of maximum weighted matchings. Another algorithm that is worth considering is the simple greedy algorithm that scans the edges in non-increasing order of length and selects an edge (u, v) if it has not previously selected an edge of the form (u, x) or (y, v) and if the collection of paths constructed so far does not include a path from v to u . On the

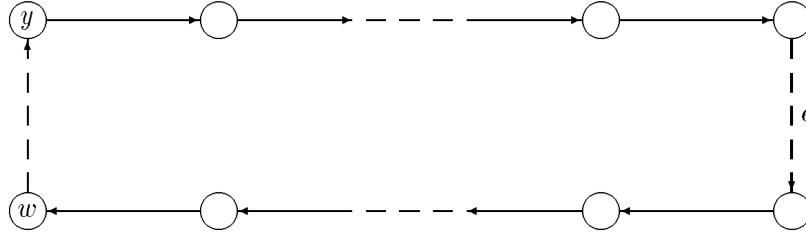


Figure 7: Illustration for Theorem 2.4

graph in Figure 6, this algorithm selects the edges $(c, f), (b, a), (e, d), (f, b), (d, c)$ in that order. The next theorem gives a worst-case bound on the performance of the greedy algorithm.

THEOREM 2.4. *If (G, ℓ) is an instance of LPP then $\lambda^*(G, \ell) \leq 3\lambda_{\text{PGREEDY}}(G, \ell)$.*

Proof. Let F be the set of edges in some optimum solution to (G, ℓ) . Let $H = \{h_1, \dots, h_s\}$ be the set of edges chosen by the greedy algorithm in the order in which they were selected (that is, h_1 was selected first, h_2 second, and so forth).

We say an edge is *permissible* at some stage of the execution of the algorithm if its selection has not been precluded by earlier selections. Define H_i to be the set of edges which are permissible just before h_i is selected, but not permissible after h_i is selected.

Let $h_i = (w, y)$ and consider the situation just before h_i is selected by the greedy algorithm. At this point, w is the last vertex of some path constructed by the algorithm and y is the first vertex of some path (one or both paths may contain just a single vertex). Let e be the edge joining the first vertex on the path containing w to the last vertex on the path containing y , as shown in Figure 7.

If e is permissible before the selection of h_i then it is a member of H_i . All other members of H_i have the form (w, z) or the form (x, y) . Note that F can contain at most one edge of the form (w, z) and at most one edge of the form (x, y) . Hence, $|F \cap H_i| \leq 3$.

Next, note that $\ell(h_i) = \max \{\ell(e) \mid e \in H_i\}$ and that (H_1, \dots, H_s) is a partition of E . Consequently, for $i \in [1, s]$, $\ell(F \cap H_i) \leq 3\ell(h_i)$ and

$$\ell(F) = \sum_{i=1}^s \ell(F \cap H_i) \leq 3 \sum_{i=1}^s \ell(h_i) = 3\ell(H) \quad \square$$

Figure 8 gives an example graph showing that the bound of Theorem 2.4 cannot be improved. (The edges not shown have length 0.) PGREEDY finds a solution of length 101, while the optimal solution has length 300. Figure 9 is a sketch of an implementation of the greedy algorithm. Upon return the mappings *left* and *right* give the left and right neighbors of each vertex in the solution path. If *left*(u) is null, then *rightend*(u) gives the vertex at the end of the path containing vertex u in the current partial solution; *leftend*(u) is similar. The running time for this implementation is $O(n^2 \log n)$.

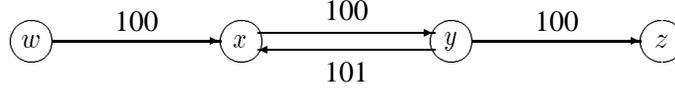


Figure 8: Worst-Case example for PGREEDY

```

function edgeset PGREEDY(digraph  $G = (V, E)$ , edgelengths  $\ell$ ,
mapping  $left, right : V \mapsto V \cup \{\text{null}\}$ )
vertex  $u, v$ ; mapping  $leftend, rightend : V \mapsto V$ ;
for  $u \in V \rightarrow$ 
   $left(u), right(u) := \text{null};$ 
   $leftend(u), rightend(u) := u;$ 
rof;
Sort  $E$  from longest to shortest;
for  $(u, v) \in E \rightarrow$ 
  if  $right(u) = \text{null}$  and  $left(v) = \text{null}$  and  $v \neq leftend(u) \rightarrow$ 
     $right(u), left(v) := v, u;$ 
     $rightend(leftend(u)) := rightend(v);$ 
     $leftend(rightend(v)) := leftend(u);$ 
  fi;
rof;
return  $S$ ;
end
  
```

Figure 9: Greedy Algorithm for LPP

3. A Greedy Algorithm for SCS. The greedy algorithm for the longest path problem can be restated for SCS as follows. Given a non-empty set of strings S , repeat the following step until S contains just one string.

Select a pair of strings $s_1, s_2 \in S$ that maximizes $\psi(s_1, s_2)$. Remove s_1 and s_2 from S , replacing them with $s_1 \circ s_2$.

We refer to this algorithm as SGREEDY. Gallant [4] claims that $\phi_{\text{SGREEDY}}(S) \leq (3/2)\phi^*(S)$. This is not in fact true, as can be seen by considering the set of strings

$$S = \{\text{abcbcbcbcb}, \text{cbcbcbcbcb}, \text{bcbcbcbcbcb}\}$$

for which $\phi_{\text{SGREEDY}}(S) = 20 > (3/2)\phi^*(S) = 19.5$. One can easily generalize this example to show that there is no constant $c < 2$ for which $\phi_{\text{SGREEDY}}(S) \leq c\phi^*(S)$. We currently do not know if there is some constant $c > 2$ for which $\phi_{\text{SGREEDY}}(S) \leq c\phi^*(S)$.

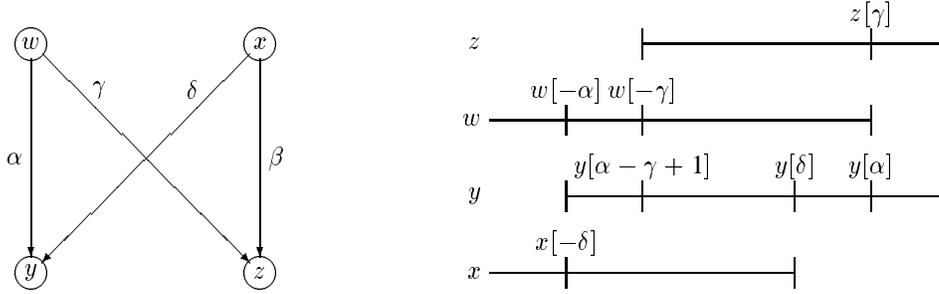


Figure 10: Illustration for Lemma 3.1

On the other hand, Theorem 2.4 allows us to conclude that $\psi^*(S) \leq 3\psi_{\text{SGREEDY}}(S)$. In fact, we can improve the constant factor to 2 by noting that the instances of LPP that arise from the transformation from SCS have a special structure which is described in the following lemma.

LEMMA 3.1. *Let S be any set of strings and let $(G, \ell) = \text{LPP}(S)$. If $\{w, x, y, z\} \subseteq V$ with $\ell(w, y) = \max \{\ell(w, y), \ell(w, z), \ell(x, y), \ell(x, z)\}$ then $\ell(w, y) + \ell(x, z) \geq \ell(w, z) + \ell(x, y)$.*

Proof. Identify w, x, y, z with the corresponding strings in S and let $\alpha = \ell(w, y)$, $\beta = \ell(x, z)$, $\gamma = \ell(w, z)$, $\delta = \ell(x, y)$. Note that if $\alpha \geq \gamma + \delta$ the result follows immediately. We will assume therefore that $\alpha \leq \gamma + \delta$. Figure 10 illustrates the situation described in the lemma.

We define some notation for designating substrings. If $s = a_1 \dots a_r$ is a string, $s[i]$ denotes the symbol a_i if $i > 0$ and a_{r+i+1} if $i < 0$. The notation $s[i, j]$ denotes the substring $s[i] \dots s[j]$.

By definition of G ,

$$w[-\alpha, -1] = y[1, \alpha] \quad (1)$$

$$w[-\gamma, -1] = z[1, \gamma] \quad (2)$$

$$x[-\delta, -1] = y[1, \delta] \quad (3)$$

Also, $w[-\gamma] = y[\alpha - \gamma + 1]$. From this we find

$$\begin{aligned} z[1, \gamma + \delta - \alpha] &= w[-\gamma, \delta - \alpha - 1] && \text{from (2), } \alpha \leq \gamma + \delta \text{ and } \alpha \geq \delta \\ &= y[\alpha - \gamma + 1, \delta] && \text{from (1)} \\ &= x[\alpha - \gamma - \delta, -1] && \text{from (3)} \end{aligned}$$

Hence, $\beta = \psi(x, z) \geq \gamma + \delta - \alpha$. \square

THEOREM 3.1. *Let S be any set of strings. $\psi^*(S) \leq 2\psi_{\text{SGREEDY}}(S)$.*

Proof. Let $(G, \ell) = \text{LPP}(S)$. Let $H = \{h_1, \dots, h_s\}$ be the set of edges chosen by the greedy algorithm in the order in which they were selected (that is, h_1 was

selected first, h_2 second, and so forth). Define $\lambda_i^* = \ell(P) - \ell(\{h_1, \dots, h_i\})$, where P is a longest path that includes $\{h_1, \dots, h_i\}$. We show that for $i \in [1, s]$, $\lambda_{i-1}^* \leq 2\ell(h_i) + \lambda_i^*$. Since, $\psi^*(S) = \lambda_0^*$, repeated expansion of this inequality implies the theorem.

Let $h_i = (w, y)$ and let $X = P - \{h_1, \dots, h_{i-1}\}$, where P is a longest path that includes $\{h_1, \dots, h_{i-1}\}$ (so $\ell(X) = \lambda_{i-1}^*$). By definition of the greedy algorithm, $\ell(w, y) = \max \{\ell(e) \mid e \in X\}$. At most three edges of X are not permissible after h_i is selected. If at most two become impermissible, then $\lambda_i^* \geq \lambda_{i-1}^* - 2\ell(w, y)$ as desired. If three edges become impermissible then one must have the form (w, z) with $z \neq y$, another the form (x, y) with $x \neq w$ and the third one, e joins the last vertex on the path containing y with the first vertex on the path containing w . This means that $X \cup \{h_1, \dots, h_{i-1}\}$ contains a path from x to z , which in turn means that (x, z) is permissible *after* (w, y) is selected. Consequently,

$$\lambda_i \geq \ell(X) - \ell(\{(x, y), (w, z), e\}) + \ell(x, z).$$

By Lemma 3.1, $\ell(w, y) + \ell(x, z) \geq \ell(w, z) + \ell(x, y)$, so

$$\ell(x, y) + \ell(w, z) + \ell(e) - \ell(x, z) \leq 2\ell(w, y)$$

which implies $\lambda_i^* \geq \lambda_{i-1}^* - 2\ell(w, y)$. \square

The bound given by Theorem 3.1 cannot be improved as can be seen by considering the set of strings mentioned at the beginning of this section. The improvement obtained for the greedy algorithm on strings raises the question of whether or not the bounds for the other approximation algorithms treated in section 2 can be improved. It turns out that they cannot. If we define

$$S = \left\{ \begin{array}{ccccc} \mathbf{a}^k \mathbf{x} \mathbf{b}^k, & \mathbf{b}^k \mathbf{x} \mathbf{c}^k, & \mathbf{c}^k \mathbf{x} \mathbf{d}^k, & \mathbf{d}^k \mathbf{x} \mathbf{e}^k, & \mathbf{e}^k \mathbf{x} \mathbf{f}^k, \\ \mathbf{b}^{k-1} \mathbf{x} \mathbf{a}^k \mathbf{x}, & \mathbf{c}^{k-1} \mathbf{x} \mathbf{b}^k \mathbf{x}, & \mathbf{d}^{k-1} \mathbf{x} \mathbf{c}^k \mathbf{x}, & \mathbf{e}^{k-1} \mathbf{x} \mathbf{d}^k \mathbf{x}, & \mathbf{f}^k \mathbf{x} \mathbf{e}^k \mathbf{x} \end{array} \right\}$$

and $(G, \ell) = \text{LPP}(S)$, we find that $\lambda^*(G, \ell) = 9k$, $\lambda_{\text{MATCH}}(G, \ell) = \lambda_{\text{DIMATCH}} = 5(k+1)$. The example can be extended to make the ratios $\lambda^*/\lambda_{\text{MATCH}}$ and $\lambda^*/\lambda_{\text{DIMATCH}}$ arbitrarily close to 2.

A naive implementation of **SGREEDY** takes at least quadratic time. A similar running time is obtained if one uses the transformation to **LPP** and then uses **PGREEDY**. A much faster running time can be obtained however by making use of an appropriate data structure. For our purposes, a *suffix tree* T is an abstract data type representing a collection of strings $S = \{s_1, \dots, s_n\}$ on which the following operations are defined.

$T.\text{SUFFIX_TREE}(S = \{s_1, \dots, s_n\})$ Initialize T to represent the strings in S . This operation may only be performed once.

$T.\text{LOOKUP}(\mathbf{integer} \ i, j)$ Returns a pair $[\ell, k]$, where ℓ is the length of the longest prefix of s_i which is also a suffix of some string in $S - \{s_j\}$ and s_k is one such string.

$T.\text{DELETE}(\mathbf{integer} \ i)$ Removes s_i from the set of strings represented.

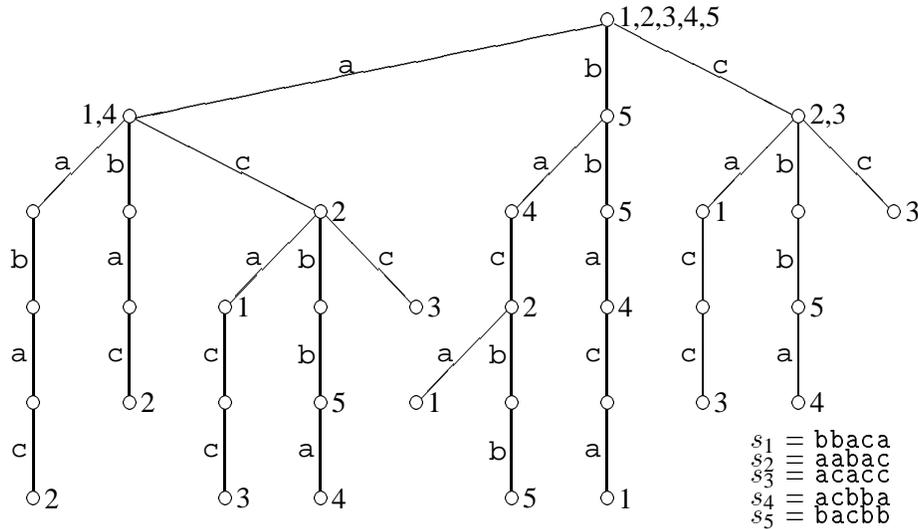


Figure 11: Example of Trie Representation for Suffix Tree

The obvious implementation of a suffix tree is a *trie* (see [1]) containing an entry for every suffix of every string in the set. An example of this representation is shown in Figure 11. The lists of integers next to some of the nodes are the indices of strings with suffixes ending at that point. This representation does not quite satisfy our needs, as the size of the trie and hence the time required to construct it is $\Omega(m^2)$ in the worst-case. A more compact representation can be obtained by labeling edges with strings rather than single characters. This allows us to eliminate many nodes with single children and results in a representation that requires $O(m)$ space and that can be constructed in $O(m)$ time, as described by McCreight [10]. See also [2, 11]. (Actually, McCreight defines a suffix tree to contain suffixes of a single string rather than a collection of strings. Our variant requires only minor modifications to McCreight’s method.) An example of this compact representation of suffix trees is shown in Figure 12.

We perform deletion in suffix trees using lazy deletion. That is, to delete a string s_i , we simply mark it deleted in an auxiliary bit vector maintained for this purpose. When a lookup operation is performed, we perform a probe in the tree to find the longest match. Let u be the node at which the probe terminates. The list of matching strings at u is scanned and any that are marked deleted are removed from the list. If this makes the list empty and u has no children, then u is removed from the tree. If no acceptable match can be found in the list, the search continues at the parent of u .

The time required for a single lookup operation may well exceed the length of the string being searched for. However, any excess time is spent deleting list entries. Since there are initially $m + n$ list entries in the whole tree, the time spent on any sequence of lookups is $O(m)$ plus the sum of the lengths of all strings being searched for.


```

procedure SGREEDY(set  $S = (s_1, \dots, s_n)$ , mapping  $left, right : [1, n] \mapsto [0, n]$ )
  integer  $i, j, \ell$ ;
  mapping  $leftend, rightend : [1, n] \mapsto [1, n]$ ;
  mapping  $key : [1, n] \mapsto \mathbf{integer}$  ;
  heap  $h$ ; suffix_tree  $T$ ;
   $T.SUFFIX\_TREE(S)$ ;
  for  $i \in [1, n] \rightarrow$ 
     $left(i), right(i) := 0$ ;
     $leftend, rightend(i) := i$ ;
     $[key(i), j] := T.LOOKUP(i, i)$ ;
     $h.INSERT(i)$ ;
  rof;
  do  $|h| > 1 \rightarrow$ 
     $i := h.FINDMAX()$ ;
     $[\ell, j] := T.LOOKUP(i, rightend(i))$ ;
    if  $\ell = key(i) \rightarrow$ 
       $left(i), right(j) := j, i$ ;
       $leftend(rightend(i)) := leftend(j)$ ;
       $rightend(leftend(j)) := rightend(j)$ ;
       $T.DELETE(j); h.DELETE(i)$ ;
    |  $\ell < key(i) \rightarrow$ 
       $key(i) := \ell; h.SIFTDOWN(i)$ ;
    fi;
  od;
end

```

Figure 13: Greedy Algorithm for SCS

As the algorithm proceeds, certain matches become unavailable and the values of $key()$ may become invalid. Consequently, whenever a string s_i is selected from h , a new lookup operation is performed in T . If the result of that operation is a match of the same length as $key(i)$, the strings are combined. If the lookup results in a shorter match, the value of $key(i)$ is changed and the position of s_i in the heap is adjusted to reflect the new value. Note that a string is deleted from the heap once it is successfully matched with another string on its left end. Similarly, a string is deleted from the suffix tree once it is matched with a string on the right.

The running time of the algorithm is dominated by the operations on the various data structures within the main loop. The number of iterations of the main loop is $O(m)$ in the worst case. Since the heap operations are $O(\log n)$ per operation, the total time spent on the heap operations is $O(m \log n)$. This can be improved to $O(m + n \log n)$ by using Fibonacci heaps [13]. Since the sequence of operations on the suffix tree is monotonic, the time needed for the suffix tree operations is $O(m)$

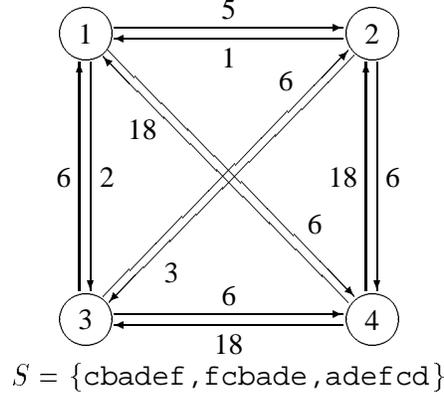


Figure 14: Example of transformation from SCS to TSP

assuming a small alphabet and $O(m \log m)$ assuming a large alphabet and the use of lexicographic splay trees. This yields a total running time of $O(m + n \log n)$ for small alphabets and $O(m \log m)$ for large alphabets.

4. SCS and the Traveling Salesman Problem. In this section we relate SCS to the path version of the *traveling salesman problem* (TSP). An instance of the traveling salesman problem is a list of cities $C = (c_1, \dots, c_n)$ with a distance $d(c_i, c_j)$ between each pair of cities. The object is to find a permutation π on $\{1, \dots, n\}$. that minimizes

$$\theta_\pi(C, d) = \sum_{i=1}^{n-1} d(c_{\pi_i}, c_{\pi_{i+1}})$$

We define $\theta^*(C, d) = \min_\pi \theta_\pi(C, d)$.

Let $S = (s_1, \dots, s_n)$ be an instance of SCS. We define $\text{TSP}(s_1, \dots, s_n)$ to be an instance (C, d) of TSP with $C = (c_1, \dots, c_n, c_{n+1})$ and

$$d(c_i, c_j) = \begin{cases} |s_i| - \psi(s_i, s_j) & 1 \leq i \leq n \quad 1 \leq j \leq n \quad i \neq j \\ |s_i| & 1 \leq i \leq n \quad j = n+1 \\ \|S\| & i = n+1 \quad 1 \leq j \leq n \end{cases}$$

An example of this transformation is given in Figure 14. Note that in general, if π satisfies $\theta_\pi(C, d) = \theta^*(C, d)$ then $\pi_{n+1} = c_{n+1}$.

THEOREM 4.1. *Let $S = (s_1, \dots, s_n)$ be an instance of SCS, $(C, d) = \text{TSP}(S)$ and let π be a permutation on $\{1, \dots, n, n+1\}$ to $(c_1, \dots, c_n, c_{n+1})$ for which $\pi_{n+1} = n+1$. Then $\theta_\pi(C, d) = \phi_{\pi'}(S)$, where π' is the restriction of π to $\{1, \dots, n\}$. In particular, $\theta^*(C, d) = \phi^*(S)$.*

Proof.

$$\begin{aligned}\theta_\pi(C, d) &= \sum_{i=1}^n d(c_{\pi_i}, c_{\pi_{i+1}}) = \left[\sum_{i=1}^{n-1} |s_{\pi_i}| - \psi(s_{\pi_i}, s_{\pi_{i+1}}) \right] + |s_{\pi_n}| \\ &= \left[\sum_{i=1}^n |s_{\pi_i}| \right] - \left[\sum_{i=1}^{n-1} \psi(s_{\pi_i}, s_{\pi_{i+1}}) \right] = \|S\| - \psi_{\pi'}(S) = \phi_{\pi'}(S)\end{aligned}$$

$\theta^*(C, d) = \phi^*(S)$ follows from the observation that any optimum solution π for (C, d) must have $\pi_{n+1} = c_{n+1}$. \square

Theorem 4.1 implies that any good approximation algorithm for this version of the traveling salesman problem is a good approximation algorithm for SCS as well. The particular instances of TSP constructed by the transformation defined above have some special properties. First, they may be asymmetric; that is $d(c_i, c_j)$ need not equal $d(c_j, c_i)$. The next theorem shows that they obey the so-called *triangle inequality*.

THEOREM 4.2. *Let $S = (s_1, \dots, s_n)$ be an instance of SCS and let $(C, d) = \text{TSP}(S)$. For all $c_i, c_j, c_k \in C$, $d(c_i, c_k) \leq d(c_i, c_j) + d(c_j, c_k)$.*

Proof. There are several cases to consider. If $i, k < j = n + 1$, $d(c_i, c_k) \leq \|S\| = d(c_j, c_k)$ and the result follows immediately. Similarly, if $i, j < k = n + 1$, $d(c_i, c_j) \leq |s_i| = d(c_i, c_k)$ and if $j, k < i = n + 1$, $d(c_i, c_j) = d(c_i, c_k)$. This leaves the case where $i, j, k < n$. For convenience, let $\alpha = d(c_i, c_j)$ and $\beta = d(c_j, c_k)$ and note that

$$s_i[\alpha + 1, |s_i|] = s_j[1, |s_i| - \alpha] \quad \text{and} \quad s_j[\beta + 1, |s_j|] = s_k[1, |s_k| - \beta]$$

Note that if $|s_i| \leq \alpha + \beta$, we are done. Therefore, assume $|s_i| > \alpha + \beta$. This implies that

$$s_i[\alpha + \beta + 1, |s_i|] = s_k[1, |s_i| - (\alpha + \beta)]$$

and that $d(c_i, c_j) \leq \alpha + \beta$. \square

There exist efficient approximation algorithms for the symmetric version of the traveling salesman problem with triangle inequality that produce solutions within a factor of $(3/2)$ of optimal. When the triangle inequality does not hold, finding good approximate solutions is as difficult as finding optimum solutions (see [7]). For the asymmetric version with triangle inequality however, little is known. There are no known approximation algorithms that are both efficient and have good worst-case performance (nor have we found any), but the approximation problem has not been shown to be hard. Consequently, the relationship between SCS and TSP has yet to yield any directly useful results. The relationship does imply some consequences if the status of either problem is resolved in the future. If good approximation algorithms are found for TSP, they may be applied to SCS. If the approximation problem for SCS is shown to be hard, then the approximation problem for TSP must be hard. It also may be that good approximation algorithms discovered for SCS, could be adapted to

TSP, although this does not necessarily follow. Finally, a proof that TSP is hard to approximate would not imply that SCS is hard to approximate, but it would imply that approximation algorithms for SCS would have to exploit special structural properties present in the instances of TSP that result from our transformation.

5. Closing Remarks. We have related the shortest common matching string problem to two other problems, using transformations that preserve solution values. These two transformations reflect different ways of viewing SCS. We have used the transformations to gain insight into the problem of approximating SCS and have discovered several algorithms that have provably good performance with respect to the overlap measure. The best of these is the string version of the greedy algorithm for which we have described an efficient implementation using suffix trees.

While we have shown that the string version of the greedy algorithm has good worst-case performance with respect to the overlap measure, we cannot determine its performance with respect to the length measure. We know that it can be off by as much as a factor of two with respect to the length measure, but we do not know if it can be worse than this. One open problem then, is to resolve this issue.

Although, we have been unable to make use of the relationship between SCS and TSP to advantage, we feel that it may yet prove useful. More generally, we think that the use of transformations that preserve solution values can be used to extend the application of known approximation algorithms to new domains. A methodical development of such transformations could provide many useful results.

Another worthwhile line of investigation for future research is to study the probable performance of the various approximation algorithms using appropriate probability models. It appears likely for example, that the directed matching algorithm for the longest path problem performs much better than its worst-case bound would indicate for a wide class of natural probability models. Similarly, one would expect the greedy algorithm to perform well in a probabilistic sense for many useful probability models.

References.

References

- [1] Aho, Alfred V., John E. Hopcroft and Jeffery D. Ullman. "Data Structures and Algorithms," Addison-Wesley, 1982.
- [2] Chen, M. T. and J. I. Seiferas. "Efficient and Elegant Subword-Tree Construction," University of Rochester, Computer Science Department, technical report TR 129, 12/83.
- [3] Gallant, John K., David Maier, James A. Storer. "On Finding Minimal Length Superstrings," *Journal of Computer and System Sciences*, vol. 20, no. 1, 50–58, 2/80.

- [4] Gallant, John K. "String Compression Algorithms," Ph.D. Dissertation, Princeton University, Department of Electrical Engineering and Computer Science, June 1982.
- [5] Gingeras, T. R., J. P. Milao, D. Sciaky and R. J. Roberts. "Computer Programs for the Assembly of DNA Sequences," *Nucleic Acids Research*, vol. 7, 1979, 529–545.
- [6] Karp, Richard M. "A Patching Algorithm for the Nonsymmetric Traveling Salesman Problem," *SIAM Journal on Computing*, vol. 8, no. 4, 561–573, 11/79.
- [7] Lawler, E. L., J. K. Lenstra and A. H. G. Rinnooy Kan (editors). "The Traveling Salesman Problem," John Wiley and Sons, Ltd, 1986.
- [8] Maier, David and James A. Storer. "A Note on the Complexity of the Superstring Problem," Princeton University Technical Report 233, Department of Electrical Engineering and Computer Science, October 1977.
- [9] Mayne, A. and E. B. James. "Information Compression by Factorising Common Strings," *The Computer Journal* 18, 1975, 157–160.
- [10] McCreight, Edward M. "A Space-Economical Suffix Tree Construction Algorithm," *Journal of the ACM*, vol. 23, 4/76, 262–272.
- [11] Rodeh, M., V. R. Pratt and S. Even. "Linear Algorithm for Data Compression via String Matching," *Journal of the ACM*, vol. 28, 1/81, 16–24.
- [12] Shapiro, M. B. "An Algorithm for Reconstructing Protein and RNA Sequences," *Journal of the ACM*, vol. 14, 1967, 720–731.
- [13] Sleator, Daniel D. and Robert E. Tarjan. "Self-Adjusting Binary Search Trees," *Journal of the ACM*, vol. 32, 7/85, 652–686.
- [14] Stefik, M. "Inferring DNA Structures From Segmentation Data," *Artificial Intelligence*, vol. 11, 1978, 85–114.
- [15] Storer, James A. and Thomas G. Szymanski. "Data Compression via Textual Substitution," *Journal of the ACM*, vol. 29, 10/82, 928–951.
- [16] Tarjan, Robert E. "Data Structures and Network Algorithms," Society for Industrial and Applied Mathematics, 1983.