# Linear Approximation of Shortest Superstrings 

AVRIM BLUM<br>Massachusetts Institute of Technology, Cambridge, Massachusetts<br>TAO JIANG<br>McMaster University, Hamilton, Ontario, Canada<br>MING LI<br>University of Waterloo, Waterloo, Ontario, Canada

JOHN TROMP
CWI, Amsterdam, The Netherlands
AND

## MIHALIS YANNAKAKIS

AT\&T Bell Laboratories, Murray Hill, New Jersey


#### Abstract

We consider the following problem: given a collection of strings $s_{1}, \ldots, s_{m}$, find the shortest string $s$ such that each $s_{i}$ appears as a substring (a consecutive block) of $s$. Although this problem is known to be NP-hard, a simple greedy procedure appears to do quite well and is routinely used in DNA sequencing and data compression practice, namely: repeatedly merge the pair of (distinct) strings with maximum overlap until only one string remains. Let $n$ denote the


[^0](c) $1994 \mathrm{ACM} 0004-5411 / 94 / 0700-0630 \$ 03.50$
length of the optimal superstring. A common conjecture states that the above greedy procedure produces a superstring of length $O(n)$ (in fact, $2 n$ ), yet the only previous nontrivial bound known for any polynomial-time algorithm is a recent $O(n \log n)$ result.

We show that the greedy algorithm does in fact achieve a constant factor approximation, proving an upper bound of $4 n$. Furthermore, we present a simple modified version of the greedy algorithm that we show produces a superstring of length at most $3 n$. We also show the superstring problem to be MAXSNP-hard, which implies that a polynomial-time approximation scheme for this problem is unlikely.
Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems-computations on discrete structures; G.2.1 [Discrete Mathematics]: Combinatorics-combinatorial algorithms
General Terms: Algorithms
Additional Key Words and Phrases: Approximation algorithm, shortest common superstring

## 1. Introduction

Given a finite set of strings, we would like to find their shortest common superstring. That is, we want the shortest possible string $s$ such that every string in the set is a substring of $s$.

The question is NP-hard [Gallant et al., 1980; Garey and Johnson, 1979]. Due to its important applications in data compression [Storer, 1988] and DNA sequencing [Lesk, 1988; Li, 1990; Peltola et al., 1983], efficient approximation algorithms for this problem are indispensable. We give an example from the DNA sequencing practice. A DNA molecule can be represented as a character string over the set of nucleotides $\{A, C, G, T\}$. Such a character string ranges from a few thousand symbols long for a simple virus to approximately $3 \times 10^{9}$ symbols for a human being. Determining this representation for different molecules, or sequencing the molecules, is a crucial step towards understanding the biological functions of the molecules. With current laboratory methods, only small fragments (chosen from unknown locations) of at most 500 bases can be sequenced at a time. Then from hundreds, thousands, sometimes millions of these fragments, a biochemist assembles the superstring representing the whole molecule. A simple greedy algorithm is routinely used [Lesk, 1988; Peltola et al., 1983] to cope with this job. This algorithm, which we call GREEDY, repeatedly merges the pair of (distinct) strings with maximum overlap until only one string remains. It has been an open question as to how well GREEDY approximates a shortest common superstring, although a common conjecture states that GREEDY produces a superstring of length at most two times optimal [Storer, 1988; Tarhio and Ukkonen, 1988; Turner, 1989].

From a different point of view, Li [1990] considered learning a superstring from randomly drawn substrings in the Valiant [1984] learning model. In a restricted sense, the shorter the superstring we obtain, the smaller the number of samples are needed to infer a superstring. Therefore, finding a good approximation bound for shortest common superstring implies efficient learnability or inferability of DNA sequences [Li, 1990]. Our linear approximation result improves Li's $O(n \log n)$ approximation by a multiplicative logarithmic factor.

Tarhio and Ukkonen [1988] and Turner [1989] established some performance guarantees for GREEDY with respect to the "compression" measure. This basically measures the number of symbols saved by GREEDY compared to plainly concatenating all the strings. It was shown that if the optimal solution
saves $l$ symbols, then GREEDY saves at least $l / 2$ symbols. But, in general, this implies no performance guarantee with respect to optimal length since in the best case this only says that GREEDY produces a superstring of length at most half the total length of all the strings.

In this paper, we show that the superstring problem can be approximated within a constant factor, and in fact that algorithm GREEDY produces a superstring of length at most $4 n$. Furthermore, we give a simple modified greedy procedure MGREEDY that also achieves a bound of $4 n$, and then present another algorithm TGREEDY, based on MGREEDY, that we show achieves $3 n$.

The rest of the paper is organized as follows: Section 2 contains notation, definitions, and some basic facts about strings. In Section 3, we describe our main algorithm MGREEDY with its proof. This proof forms the basis of the analysis in the next two sections. MGREEDY is improved to TGREEDY in Section 4. We finally give the $4 n$ bound for GREEDY in Section 5. Section 6 presents a simple comparison of the performance of these algorithms. In Section 7, we show that the superstring problem is MAX SNP-hard that implies that there is unlikely to exist a polynomial-time approximation scheme for the superstring problem.

## 2. Preliminaries

Let $S=\left\{s_{1}, \ldots, s_{m}\right\}$ be a set of strings over some alphabet $\Sigma$. Without loss of generality, we assume that the set $S$ is "substring-free" in that no string $s_{i} \in S$ is a substring of any other $s_{j} \in S$. A common superstring of $S$ is a string $s$ such that each $s_{i}$ in $S$ is a substring of $s$. That is, for each $s_{i}$, the string $s$ can be written as $u_{i} s_{i} v_{i}$ for some $u_{i}$ and $v_{i}$. In this paper, we use $n$ and $\operatorname{OPT}(S)$ interchangeably for the length of the shortest common superstring for $S$. Our goal is to find a superstring for $S$ whose length is as close to OPT $(S)$ as possible.

Example. Assume we want to find the shortest common superstring of all words in the following sentence: "Alf ate half lethal alpha alfalfa". The word "alf" is a substring of both "half" and "alfalfa", so we can immediately eliminate it. Our set of words is now $S_{0}=\{$ ate, half, lethal, alpha, alfalfa\}. A trivial superstring is "atehalflethalalphaalfalfa" of length 25 , which is simply the concatenation of all substrings. A shortest common superstring is "lethalphalfalfate", of length 17 , saving 8 characters over the previous one (a compression of 8). Looking at what GREEDY would make of this example, we see that it would start out with the largest overlaps from "lethal" to "half" to "alfalfa" producing "lethalfalfa". It then has 3 choices of single character overlap, two of which lead to another shortest superstring "lethalfalfalphate", and one of which is lethal in the sense of giving a superstring that is one character longer. In fact, it is easy to give an example where GREEDY outputs a string almost twice as long as the optimal one, for instance on input $\left\{c(a b)^{k},(b a)^{k},(a b)^{k} c\right\}$.

For two strings $s$ and $t$, not necessarily distinct, let $v$ be the longest string such that $s=u v$ and $t=v w$ for some non-empty strings $u$ and $w$. We call $|v|$ the (amount of) overlap between $s$ and $t$, and denote it as ov( $s, t$ ). Furthermore, $u$ is called the prefix of $s$ with respect to $t$, and is denoted $\operatorname{pref}(s, t)$. Finally, we call $|\operatorname{pref}(s, t)|=|u|$ the distance from $s$ to $t$, and denote it as
$d(s, t)$. So, the string $u v w=\operatorname{pref}(s, t) t$, of length $d(s, t)+|t|=|s|+|t|-$ $o v(s, t)$ is the shortest superstring of $s$ and $t$ in which $s$ appears (strictly) before $t$, and is also called the merge of $s$ and $t$. For $s_{i}, s_{j} \in S$, we will abbreviate $\operatorname{pref}\left(s_{i}, s_{j}\right)$ to simply $\operatorname{pref}(i, j)$, and $d\left(s_{i}, s_{j}\right)$ and $o v\left(s_{i}, s_{j}\right)$ to $d(i, j)$ and ov $(i, j)$ respectively. The overlap between a string and itself is called a self-overlap. As an example of self-overlap, we have for the string $s=$ undergrounder an overlap of $o v(s, s)=5$. Also, $\operatorname{pref}(s, s)=$ undergro and $d(s, s)=8$. The string $s=$ alfalfa, for which $o v(s, s)=4$, shows that the overlap is not limited to half the total string length.

Given a list of strings $s_{i_{1}}, s_{i_{2}}, \ldots, s_{i_{r}}$, we define the superstring $s=\left\langle s_{i_{1}}, \ldots, s_{i_{1}}\right\rangle$ to be the string $\operatorname{pref}\left(i_{1}, i_{2}\right) \operatorname{pref}\left(i_{2}, i_{3}\right) \cdots \operatorname{pref}\left(i_{r-1}, i_{r}\right) s_{i,}$. That is, $s$ is the shortest string such that $s_{i}, s_{i}, \ldots, s_{i,}$ appear in order in that string. For a superstring of a substring-free set, this order is well-defined, since substrings cannot "start" or "end" at the same position, and if substring $s_{j}$ starts before $s_{k}$, then $s_{j}$ must also end before $s_{k}$. Define first $(s)=s_{i}$ and last $(s)=s_{i_{i}}$. In each iteration of GREEDY the following invariant holds:

Clalm 1. For two distinct strings $s$ and $t$ in GREEDY's set of strings, neither first $(s)$ nor last ( $s$ ) is a substring of $t$.

Proof. Initially, first $(s)=\operatorname{last}(s)=s$ for all strings, so the claim follows from the fact that $S$ is substring-free. Suppose that the invariant is invalidated by a merge of two strings $t_{1}$ and $t_{2}$ into a string $t=\left\langle t_{1}, t_{2}\right\rangle$ that has, say, first $(s)$ as a substring. Let $t=u$ first $(s) v$. Since first $(s)$ is not a substring of either $t_{1}$ or $t_{2}$, it must properly "contain" the piece of overlap between $t_{1}$ and $t_{2}$, that is, $\mid$ first $(s) \mid>o v\left(t_{1}, t_{2}\right)$ and $|u|<d\left(t_{1}, t_{2}\right)$. Hence, ov $\left(t_{1}, s\right)>o{ }^{\prime}\left(t_{1}, t_{2}\right)$; a contradiction.

So when GREEDY (or its variation MGREEDY that we introduce later) chooses $s$ and $t$ as having the maximum overlap, then this overlap ov $(s, t)$ in fact equals ou(last $(s)$, first $(t)$ ), and as a result, the merge of $s$ and $t$ is $\langle\operatorname{first}(s), \ldots, \operatorname{last}(s)$, first $(t), \ldots, \operatorname{last}(t)\rangle$. We can therefore say that GREEDY orders the substrings, by finding the shortest superstring in which the substrings appear in that order.

We can rephrase the above in terms of permutations. For a permutation $\pi$ on the set $\{1, \ldots, m\}$, let $S_{\pi}=\left\langle s_{\pi(1)}, \ldots, s_{\pi(m)}\right\rangle$. In a shortest superstring for $S$, the substrings appear in some total order, say $s_{\pi(1)}, \ldots, s_{\pi(m)}$, hence it must equal $S_{\pi}$.

We consider a traveling salesman problem on a weighted directed complete graph $G_{S}$ derived from $S$ and show that one can achieve a factor of 4 approximation for TSP on that graph, yielding a factor of 4 approximation for the shortest-common-superstring problem. Graph $G_{S}=(V, E, d)$ has $m$ vertices $V=\{1, \ldots, m\}$, and $m^{2}$ edges $E=\{(i, j): 1 \leq i, j \leq m\}$. Here we take as weight function the distance $d($,$) : edge ( i, j$ ) has weight $d(i, j)=d\left(s_{i}, s_{j}\right)$, to obtain the distance graph. This graph is similar to one considered by Turner [1989] in the end of his paper. Later, we take the overlap ov(,) as the weight function to obtain the overlap graph. We call $s_{i}$ the string associated with vertex $i$, and let $\operatorname{pref}(i, j)=\operatorname{pref}\left(s_{i}, s_{j}\right)$ be the string associated with edge $(i, j)$.

As examples, we draw in Figure 1 the overlap graph and the distance graph for our previous example $S_{0}=\{$ ate, half, lethal, alpha, alfalfa $\}$. All edges not shown have overlap 0 . Note that the sum of the distance and overlap weights on an edge $(i, j)$ is the length of the string $s_{i}$.


Fig. 1. The overlap and distance graphs.

Notice now that $\operatorname{TSP}\left(G_{S}\right) \leq \operatorname{OPT}(S)-\operatorname{ov}(\operatorname{last}(s)$, first $(s)) \leq \operatorname{OPT}(S)$, where $\operatorname{TSP}\left(G_{S}\right)$ is the cost of the minimum weight Hamiltonian cycle on $G_{S}$. The reason is that turning any superstring into a Hamiltonian cycle by overlapping its last and first substring saves on cost by charging last $(s)$ for only $d(\operatorname{last}(s)$, first $(s))$ instead of its full length.

We now define some notation for dealing with directed cycles in $G_{S}$. Call two strings $s, t$ equivalent, $s \equiv t$, if they are cyclic shifts of each other, that is, if there are strings $u, v$ such that $s=u v$ and $t=v u$. If $c$ is a directed cycle in $G_{S}$ with vertices $i_{0}, \ldots, i_{r-1}$ in order around $c$, we define strings $(c)$ to be the equivalence class $\left[\operatorname{pref}\left(i_{0}, i_{1}\right) \operatorname{pref}\left(i_{1}, i_{2}\right) \cdots \operatorname{pref}\left(i_{r-1}, i_{0}\right)\right]$ and $\operatorname{strings}\left(c, i_{k}\right)$ the rotation starting with $\operatorname{pref}\left(i_{k}, i_{k+1}\right)$, that is, the string $\operatorname{pref}\left(i_{k}, i_{k+1}\right) \cdots$ $\operatorname{pref}\left(i_{k-1}, i_{k}\right)$, where subscript arithmetic is modulo $r$. Let us say that an equivalence class $[s$ ] has periodicity $k(k>0)$, if $s$ is invariant under a rotation by $k$ characters ( $s=u v=v u,|u|=k$ ). Obviously, $[s]$ has periodicity $|s|$. A moment's reflection shows that the minimum periodicity of [ $s$ ] must equal the number of distinct rotations of $s$. That is the size of the equivalence class and denoted by $\operatorname{card}([s])$. Furthermore, it is easily proven that if $[s]$ has periodicities $a$ and $b$, then it has periodicity $\operatorname{gcd}(a, b)$ as well. (See, e.g., Fine and Wilf [1965].) It follows that all periodicities are a multiple of the minimum one. In particular, we have that $|s|$ is a multiple of $\operatorname{card}([s])$.

In general, we denote a cycle $c$ with vertices $i_{1}, \ldots, i_{r}$ in the order by " $i_{1} \rightarrow \cdots \rightarrow i_{r} \rightarrow i_{1}$." Also, let $w(c)$, the weight of cycle $c$, equal $|s|, s \in$ strings ( $c$ ). For convenience, we say that $s_{j}$ is in $c$, or " $s_{j} \in c$ " if $j$ is a vertex of the cycle $c$.

Now, a few preliminary facts about cycles in $G_{S}$. Let $c=i_{0} \rightarrow \cdots \rightarrow i_{r-1} \rightarrow$ $i_{0}$ and $c^{\prime}$ be cycles in $G_{s}$. For any string $s, s^{k}$ denotes the string consisting of $k$ copies of $s$ concatenated together.

Claim 2. Each string $s_{i_{j}}$ in $c$ is a substring of $s^{k}$ for all $s \in \operatorname{strings(c)~and~}$ sufficiently large $k$.

Proof. By induction, $s_{i_{j}}$ is a prefix of $\operatorname{pref}\left(i_{j}, i_{j+1}\right) \cdots \operatorname{pref}\left(i_{j+1-1}, i_{j+l}\right) s_{i_{j+1}}$ for any $l \geq 0$ (addition modulo $r$ ). Taking $k=\left\lceil\left|s_{i_{j}}\right| / w(c)\right\rceil$ and $l=k r$, we get that $s_{i_{j}}$ is a prefix of $\operatorname{pref}\left(i_{j}, i_{j+1}\right) \cdots \operatorname{pref}\left(i_{j+k r-1}, i_{j+k r}\right)=\operatorname{strings}\left(c, i_{j}\right)^{k}$, which itself is a substring of $s^{k+1}$ for any $s \in \operatorname{strings}(c)$.

Claim 3. If each of $\left\{s_{j}, \ldots, s_{j}\right\}$ is a substring of $s^{k}$ for some string $s \in$ strings( $c$ ) and sufficiently large $k$, then there exists a cycle of weight $|s|=w(c)$ containing all these strings.

Proof. In a (infinite) repetition of $s$, every string $s_{i}$ appears as a substring after every $|s|$ characters. This naturally defines a circular ordering of the strings $\left\{s_{j_{1}}, \ldots, s_{j_{r}}\right\}$ and the strings in $c$ whose successive distances sum to $|s|$.

Claim 4. The superstring $\left\langle s_{i_{0}}, \ldots, s_{i_{r-1}}\right\rangle$ is a substring of strings $\left(c, i_{0}\right) s_{i_{0}}$.
Proof. String $\left\langle s_{i_{0}}, \ldots, s_{i_{r-1}}\right\rangle$ is clearly a substring of $\left\langle s_{i_{0}}, \ldots, s_{i_{r-1}}, s_{i_{0}}\right\rangle$, which by definition equals $\operatorname{pref}\left(i_{0}, i_{1}\right) \cdots \operatorname{pref}\left(i_{r-1}, i_{0}\right) s_{i_{0}}=\operatorname{strings}\left(c, i_{0}\right) s_{i_{0}}$.

Claim 5. If strings $\left(c^{\prime}\right)=\operatorname{strings}(c)$, then there exists a third cycle $\tilde{c}$ with weight $w(c)$ containing all vertices in $c$ and all those in $c^{\prime}$.

Proof. Follows from claims 2 and 3.
Claim 6. There exists a cycle $\bar{c}$ of weight card(strings(c)) containing all vertices in $c$.

Proof. Let $u$ be the prefix of length card(strings( $c)$ ) of some strings $s \in \operatorname{string} s(c)$. By our periodicity arguments, $|u|$ divides $|s|=w(c)$, and $s=u^{j}$ where $j=w(c) /|u|$. It follows that every string in strings $(c)=[s]$ is a substring of $u^{j+1}$. Now use Claim 3 for strings $(c)$ and $u$.

The following lemma has been proved in Tarhio and Ukkonen [1988] and Turner [1989]. Figure 2 gives a graphical interpretation of it. In the figure, the vertical bars surround pieces of string that match, showing a possible overlap between $v^{-}$and $u^{+}$, giving an upper bound on $d\left(v^{-}, u^{+}\right)$.

Lemma 7. Let $u, u^{+}, v^{-}, v$ be strings, not necessarily different, such that $o v(u, v) \geq \max \left\{o v\left(u, u^{+}\right), o v\left(v^{-}, v\right)\right\}$. Then, ov $(u, v)+o v\left(v^{-}, u^{+}\right) \geq o v\left(u, u^{+}\right)$ $+o v\left(v^{-}, v\right)$, and $d(u, v)+d\left(v^{-}, u^{+}\right) \leq d\left(u, u^{+}\right)+d\left(v^{-}, v\right)$.

That is, given the choice of merging $u$ to $u^{+}$and $v^{-}$to $v$ or instead merging $u$ to $v$ and $v^{-}$to $u^{+}$, the best choice is that which contains the pair of largest overlap. The conditions in the above lemma are also known as "Monge conditions" in the context of transportation problems [Alon et al., 1989; Barnes and Hoffman, 1985; Hoffman, 1963]. In this sense, the lemma follows from the observation that optimal shipping routes do not intersect. In the string context, we are transporting "items" from the ends of substrings to the fronts of substrings.

## 3. A $4 \cdot \operatorname{OPT}(S)$ Bound for a Modified Greedy Algorithm

Let $S$ be a set of strings and $G_{S}$ the associated graph. Now, although finding a minimum weight Hamiltonian cycle in a weighted directed graph is in general a hard problem, there is a polynomial-time algorithm for a similar problem known as the assignment problem [Papadimitriou and Steiglitz, 1982]. Here, the goal is simply to find a decomposition of the graph into cycles such that each vertex is in exactly one cycle and the total weight of the cycles is minimized. Let $\mathrm{CYC}\left(G_{S}\right)$ be the weight of the minimum assignment on graph $G_{S}$, so $\operatorname{CYC}\left(G_{S}\right) \leq \operatorname{TSP}\left(G_{S}\right) \leq \operatorname{OPT}(S)$.

The proof that a modified greedy algorithm MGREEDY finds a superstring of length at most $4 \cdot \mathrm{OPT}(S)$ proceeds in two stages. We first show that an algorithm that finds an optimal assignment on $G_{S}$, then opens each cycle into a single string, and finally concatenates all such strings together has a


Fig. 2. Strings and overlaps.
performance ratio of at most 4. We then show (Theorem 10) that in fact, for these particular graphs, a greedy strategy can be used to find optimal assignments. This result can also be found (in a somewhat different form) as Theorem 1 in Hoffman [1963].

Consider the following algorithm for finding a superstring of the strings in $S$.

## Algorithm Concat-Cycles

(1) On input $S$, create graph $G_{S}$ and find a minimum weight assignment $C$ on $G_{S}$. Let $C$ be the collection of cycles $\left\{c_{1}, \ldots, c_{p}\right\}$.
(2) For each cycle $c_{i}=i_{1} \rightarrow \cdots \rightarrow i_{r} \rightarrow i_{1}$, let $\tilde{s}_{i}=\left\langle s_{i_{1}}, \ldots, s_{i_{r}}\right\rangle$ be the string obtained by opening $c_{i}$, where $i_{1}$ is arbitrarily chosen. The string $\bar{s}_{i}$ has length at most $w\left(c_{i}\right)+\left|s_{i}\right|$ by Claim 4 .
(3) Concatenate together the strings $\tilde{s}_{i}$ and produce the resulting string $\tilde{s}$ as output.

Theorem 8. Algorithm Concat-Cycles produces a string of length at most 4 - OPT(S).

Before proving Theorem 8, we first need a preliminary lemma giving an upper bound on the amount of overlap possible between strings in different cycles of $C$. The lemma is also implied by the results in Fine and Wilf [1965].

Lemma 9. Let $c$ and $c^{\prime}$ be two cycles in a minimum weight assignment $C$ with $s \in c$ and $s^{\prime} \in c^{\prime}$. Then, the overlap between $s$ and $s^{\prime}$ is less than $w(c)+w\left(c^{\prime}\right)$.

Proof. Let $x=\operatorname{strings}(c)$ and $x^{\prime}=\operatorname{strings}\left(c^{\prime}\right)$. Since $C$ is a minimum weight assignment, we know $x \neq x^{\prime}$. Otherwise, by Claim 5 , we could find a lighter assignment by combining the cycles $c$ and $c^{\prime}$. In addition, by Claim 6, $w(c) \leq \operatorname{card}(x)$.

Suppose that $s$ and $s^{\prime}$ overlap in a string $u$ with $|u| \geq w(c)+w\left(c^{\prime}\right)$. Denote the substring of $u$ starting at the $i$ th symbol and ending at the $j$ th as $u_{i, j}$. Since, by Claim 2, $s=t^{k}$ for some $t \in x$ and large enough $k$ and $s^{\prime}=t^{\prime k^{\prime}}$ for some $t^{\prime} \in x^{\prime}$ and large enough $k^{\prime}$, we have that $x=\left[u_{1, w(c)}\right]$ and $x^{\prime}=\left[u_{1, w\left(c^{\prime}\right)}\right]$. From $x \neq x^{\prime}$, we conclude that $w(c) \neq w\left(c^{\prime}\right)$; assume without loss of generality that $w(c)>w\left(c^{\prime}\right)$. Then

$$
u_{1, w(c)}=u_{1+w\left(c^{\prime}\right), w(c)+w\left(c^{\prime}\right)}=u_{1+w\left(c^{\prime}\right), w(c)} u_{w(c)+1, w(c)+w\left(c^{\prime}\right)}=u_{1+w\left(c^{\prime}\right), w(c)} u_{1, w\left(c^{\prime}\right)} .
$$

This shows that $x$ has periodicity $w\left(c^{\prime}\right)<w(c) \leq \operatorname{card}(x)$, which contradicts the fact that $\operatorname{card}(x)$ is the minimum periodicity.

Proof of Theorem 8. Since $C=\left\{c_{1}, \ldots, c_{p}\right\}$ is an optimal assignment, $\operatorname{CYC}\left(G_{S}\right)=\sum_{i=1}^{P} w\left(c_{i}\right) \leq \operatorname{OPT}(S)$. A second lower bound on $\operatorname{OPT}(S)$ can be determined as follows: For each cycle $c_{i}$, let $w_{i}=w\left(c_{i}\right)$ and $l_{i}$ denote the length of the longest string in $c_{i}$. By Lemma 9, if we consider the longest string in each cycle and merge them together optimally, the total amount of overlap will be at most $2 \sum_{i=1}^{p} w_{i}$. So the resulting string will have length at least $\sum_{i=1}^{p} l_{i}-2 w_{i}$. Thus, $\operatorname{OPT}(S) \geq \max \left(\sum_{i=1}^{p} w_{i}, \sum_{i=1}^{p} l_{i}-2 w_{i}\right)$.

The output string $\tilde{s}$ of algorithm Concat-Cycles has length at most $\sum_{i=1}^{p} l_{i}+w_{i}$ (Claim 4). So,

$$
\begin{aligned}
|\tilde{s}| & \leq \sum_{i=1}^{p} l_{i}+w_{i} \\
& =\sum_{i=1}^{p} l_{i}-2 w_{i}+\sum_{i=1}^{p} 3 w_{i} \\
& \leq \operatorname{OPT}(S)+3 \cdot \operatorname{OPT}(S) \\
& =4 \cdot \operatorname{OPT}(S) .
\end{aligned}
$$

We are now ready to present the algorithm MGREEDY, and show that it in fact mimics algorithm Concat-Cycles.

## Algorithm MGREEDY

(1) Let $S$ be the input set of strings and $T$ be empty.
(2) While $S$ is non-empty, do the following: Choose $s, t \in S$ (not necessarily distinct) such that $o l(s, t)$ is maximized, breaking ties arbitrarily. If $s \neq t$, then remove $s$ and $t$ from $S$ and replace them with the merged string $\langle s, t\rangle$. If $s=t$, then just remove $s$ from $S$ and add it to $T$.
(3) When $S$ is empty, output the concatenation of the strings in $T$.

We can look at MGREEDY as choosing edges in the overlap graph ( $V=$ $S, E=V \times V$,ou $()$,$) . When MGREEDY chooses strings s$ and $t$ as having the maximum overlap (where $t$ may equal $s$ ), it chooses the directed edge from last (s) to first $(t)$ (see Claim 1). Thus, MGREEDY constructs/joins paths, and closes them into cycles, to end up with a collection of disjoint cycles $M \subset E$ that cover the vertices of $G_{S}$. We will call $M$ the assignment created by MGREEDY. Now think of MGREEDY as taking a list of all the edges sorted in the decreasing order of their overlaps (resolving ties in some definite way), and going down the list deciding for each edge whether to include it or not. Let us say that an edge e dominates another edge $f$ if $e$ precedes $f$ in this list and shares its head (or tail) with the head (or tail, respectively) of $f$. By the definition of MGREEDY, it includes an edge $f$ if and only if it has not yet included an edge dominating $f$.

THEOREM 10. The assignment created by algorithm MGREEDY is an optimal assignment.

Proof. Note that the overlap weight of an assignment and its distance weight add up to the total length of all strings. Accordingly, an assignment is optimal (i.e., has minimum total weight in the distance graph) if and only if it has maximum total overlap. Among the maximum overlap assignments, let $N$ be one that has the maximum number of edges in common with $M$. We shall show that $M=N$.

Suppose this is not the case, and let $e$ be the edge of maximum overlap in the symmetric difference of $M$ and $N$, with ties broken the same way as by MGREEDY. Suppose first that this edge is in $N \backslash M$. Since MGREEDY did not include $e$, it must have included another adjacent edge $f$ that dominates $e$. Edge $f$ cannot be in $N$ (since $N$ is an assignment), therefore $f$ is in $M \backslash N$, contradicting our choice of the edge $e$. Suppose that $e=k \rightarrow j$ is in $M \backslash N$. The two $N$ edges $i \rightarrow j$ and $k \rightarrow l$ that share head and tail with $e$ are not in $M$,
and thus are dominated by $e$. Since $o v(k, j) \geq \max \{o v(i, j), o v(k, l)\}$, by Lemma 7 , $o v(i, j)+o v(k, l) \leq o v(k, j)+o v(i, l)$. Thus, replacing in $N$, these two edges with $e=k \rightarrow j$ and $i \rightarrow l$ would yield an assignment $N^{\prime}$ that has more edges in common with $M$ and has no less overlap than $N$. This would contradict our choice of $N$.

Since algorithm MGREEDY finds an optimal assignment, the string it produces is no longer than the string produced by algorithm Concat-Cycles. (In fact, it could be shorter since it breaks each cycle in the optimum position.)

## 4. Improving to $3 \cdot O P T(S)$

Recall that in the last step of algorithm MGREEDY, we simply concatenate all the strings in set $T$ without any compression. Intuitively, if we instead try to overlap the strings in $T$, we might be able to achieve a bound better than $4 \cdot \operatorname{OPT}(S)$. Let TGREEDY denote the algorithm that operates in the same way as MGREEDY except that in the last step, it merges the strings in $T$ by running GREEDY on them. We can show that TGREEDY indeed achieves a better bound: it produces a superstring of length at most $3 \cdot \mathrm{OPT}(S)$.

Theorem 11. Algorithm TGREEDY produces a superstring of length at most $3 \cdot O P T(S)$.

Proof. Let $S=\left\{s_{1}, \ldots, s_{m}\right\}$ be a set of strings and $s$ be the superstring obtained by TGREEDY on $S$. Let $n=\operatorname{OPT}(S)$ be the length of a shortest superstring of $S$. We show that $|s| \leq 3 n$.

Let $T$ be the set of all "self-overlapping" strings obtained by MGREEDY on $S$ and $C$ be the assignment created by MGREEDY. For each $x \in T$, let $c_{x}$ denote the cycle in $C$ corresponding to string $x$, and let $w_{x}=w\left(c_{x}\right)$ be its weight. For any set $R$ of strings, define $\|R\|=\sum_{x \in R}|x|$ to be the total length of the strings in set $R$. Also let $w=\sum_{x \in T} w_{x}$. Since $\operatorname{CYC}\left(G_{S}\right) \leq \operatorname{TSP}\left(G_{S}\right) \leq$ OPT( $S$ ), we have $w \leq n$.

By Lemma 9, the compression achieved in a shortest superstring of $T$ is less than $2 w$, that is, $\|T\|-n_{T} \leq 2 w$. By the results in Tarhio and Ukkonen [1983] and Turner [1989], we know that the compression achieved by GREEDY on set $T$ is at least half the compression achieved in any superstring of $T$. That is,

$$
\|T\|-|s| \geq \frac{\|T\|-n_{T}}{2}=\|T\|-n_{T}-\frac{\|T\|-n_{T}}{2} \geq\|T\|-n_{T}-w .
$$

So, $|s| \leq n_{T}+w$.
For each $x \in T$, let $s_{i_{x}}$ be the string in cycle $c_{x}$ that is a prefix of $x$. Let $S^{\prime}=\left\{s_{i_{x}} \mid x \in T\right\}, \quad n^{\prime}=\operatorname{OPT}\left(S^{\prime}\right), \quad S^{\prime \prime}=\left\{\operatorname{strings}\left(c_{x}, i_{x}\right) s_{i_{x}} \mid x \in T\right\}, \quad$ and $\quad n^{\prime \prime}=$ OPT ( $S^{\prime \prime}$ ).

By Claim 4, a superstring for $S^{\prime \prime}$ is also a superstring for $T$, so $n_{T} \leq n^{\prime \prime}$, where $n_{T}=\operatorname{OPT}(T)$. For any permutation $\pi$ on $T$, we have $\left|S_{\pi}^{\prime \prime}\right| \leq\left|S_{\pi}^{\prime}\right|+$ $\sum_{x \in T} w_{x}$, so $n^{\prime \prime} \leq n^{\prime}+w$, where $S_{\pi}^{\prime}$ and $S_{\pi}^{\prime \prime}$ are the superstrings obtained by overlapping the members of $S^{\prime}$ and $S^{\prime \prime}$, respectively, in the order given by $\pi$. Observe that $S^{\prime} \subseteq S$ implies $n^{\prime} \leq n$. Summing up, we get

$$
n_{T} \leq n^{\prime \prime} \leq n^{\prime}+w \leq n+w
$$

Combined with $|s| \leq n_{T}+w$, this gives $|s| \leq n+2 w \leq 3 n$.

## 5. GREEDY Achieves Linear Approximation

One would expect that an analysis similar to that of MGREEDY would also work for the original GREEDY. This turns out not to be the case. The analysis of GREEDY is severely complicated by the fact that it continues processing the "self-overlapping" strings. MGREEDY was especially designed to avoid these complications, by separating such strings. Let $G R E E D Y(S)$ denote the length of the superstring produced by GREEDY on a set $S$. It is tempting to claim that

$$
G R E E D Y(S \cup\{s\}) \leq G R E E D Y(S)+|s|
$$

If this were true, a simple argument would extend the $4 \cdot \mathrm{OPT}(S)$ result for MGREEDY to GREEDY. But the following counterexample disproves this seemingly innocent claim. Let

$$
S=\left\{c a^{m}, a^{m+1} c^{m}, c^{m} b^{m+1}, b^{m} c\right\}, \quad s=b^{m+1} a^{m+1} .
$$

Now

$$
\operatorname{GREEDY}(S)=\left|c a^{m+1} c^{m} b^{m+1} c\right|=3 m+4,
$$

whereas

$$
\begin{aligned}
\operatorname{GREEDY}(S \cup\{s\}) & =\left|b^{m} c^{m} b^{m+1} a^{m+1} c^{m} a^{m}\right| \\
& =6 m+2>(3 m+4)+(2 m+2)
\end{aligned}
$$

With a more complicated analysis, we nevertheless show that
THEOREM 12. GREEDY produces a string of length at most $4 \cdot O P T(S)$.
Before proving the theorem formally, we give a sketch of the basic idea behind the proof. If we want to relate the merges done by GREEDY to an optimal assignment, we have to keep track of what happens when GREEDY violates the maximum overlap principle, that is, when some self-overlap is better than the overlap in GREEDY's merge. One thing to try is to charge GREEDY some extra cost that reflects that an optimal assignment on the new set of strings (with GREEDY's merge) may be somewhat longer than the optimal assignment on the former set (in which the self-overlapping string would form a cycle). If we could just bound these extra costs, then we would have a bound for GREEDY. Unfortunately, this approach fails because the self-overlapping string may be merged by GREEDY into a larger string which itself becomes self-overlapping, and this nesting could go arbitrarily deep. Our proof concentrates on the innermost self-overlapping strings only. These so called culprits form a linear order in the final superstring. We avoid the complications of higher level self-overlaps by splitting the analysis in two parts. In one part, we ignore all the original substrings that connect first to the right of a culprit. In the other part, we ignore all the original substrings that connect first to the left of a culprit. In each case, it becomes possible to bound the extra cost. This method yields a bound of $7 \cdot \mathrm{OPT}(S)$. By combining the two analyses in a more clever way, we can even eliminate the effect of the extra costs and obtain the same $4 \cdot \operatorname{OPT}(S)$ bound as we found for MGREEDY. A detailed formal proof follows:

Proof of Theorem 12. We need some notions and lemmas. Think of both GREEDY and MGREEDY as taking a list of all edges sorted by overlap, and
going down the list deciding for each edge whether to include it or not. Call an edge better (worse) if it appears before (after) another in this list. Better edges have at least the overlap of worse ones. Recall that an edge dominates another iff it is better and shares its head or tail with the other one.

At the end, GREEDY has formed a Hamiltonian path

$$
s_{1} \rightarrow s_{2} \rightarrow \cdots \rightarrow s_{m}
$$

of "greedy" edges. (Without loss of generality, the strings are renumbered to reflect their order in the superstring produced by GREEDY.) For convenience we usually abbreviate $s_{i}$ to $i$. GREEDY does not include an edge $f$ iff
(1) $f$ is dominated by an already chosen edge $e$, or
(2) $f$ is not dominated but it would form a cycle.

Let us call the latter "bad back edges"; a bad back edge $f=j \rightarrow i$ necessarily has $i \leq j$. Each bad back edge $f=j \rightarrow i$ corresponds to a string $\left\langle s_{i}, s_{i+1}, \ldots, s_{j}\right\rangle$ that, at some point in the execution of GREEDY, has more (self) overlap than the pair that is merged. When GREEDY considers $f$, it has already chosen all (better) edges on the greedy path from $i$ to $j$, but not yet the (worse) edges $i-1 \rightarrow i$ and $j \rightarrow j+1$. The bad back edge $f$ is said to span the closed interval $I_{f}=[i, j]$. The above observations provide a proof of the following lemma.

Lemma 13. Let e and fo two bad back edges. The closed intervals $I_{e}$ and $I_{f}$ are either disjoint, or one contains the other. If $I_{e} \supset I_{f}$, then $e$ is worse than $f$ (thus, ol $(e) \leq o v(f)$ ).

Thus, the intervals of the bad back edges are nested and bad back edges do not cross each other. Culprits are the minimal (innermost) such intervals. Each culprit $[i, j]$ corresponds to a culprit string $\left\langle s_{i}, s_{i+1}, \ldots, s_{j}\right\rangle$. Note that, because of the minimality of the culprits, if $f=j \rightarrow i$ is the back edge of a culprit $[i, j]$, and $e$ is another bad back edge that shares head or tail with $f$, then $I_{c} \supset I_{f}$, and therefore $f$ dominates $e$.

Call the worst edge between every two successive culprits on the greedy path a weak link. Note that weak links are also worse than all edges in the two adjacent culprits as well as their back edges. If we remove all the weak links, the greedy path is partitioned into a set of paths, called blocks. Every block consists of a nonempty culprit as the middle segment, and (possibly empty) left and right extensions. The set of strings (nodes) $S$ is thus partitioned into three sets $S_{l}, S_{m}, S_{r}$ of left, middle, and right strings. The example in Figure 3 has seven substrings, of which 2 by itself and the merge of 4,5 , and 6 form the culprits (indicated by thicker lines). Bad back edges are $2 \rightarrow 2,6 \rightarrow 4$, and $6 \rightarrow 1$. The weak link $3 \rightarrow 4$ is the worst edge between culprits [2] and [4, 5, 6]. The blocks in this example are thus $[1,2,3]$ and $[4,5,6,7]$, and we have $S_{l}=\{1\}$, $S_{m}=\{2,4,5,6\}, S_{r}=\{3,7\}$.

The following lemma shows that a bad back edge must be from a middle or right node to a middle or left node.

Lemma 14. Let $f=j \rightarrow i$ be a bad back edge. Node $i$ is either a left node or the first node of a culprit. Node $j$ is either a right node or the last node of a culprit.
Proof. Let $c=[k, l]$ be the leftmost culprit in $I_{f}$. Now either $i=k$ is the first node of $c$, or $i<k$ is in the left extension of $c$, or $i<k$ is in the right


Fig. 3. Culprits and weak links in Greedy merge path.
extension of the culprit $c^{\prime}$ to the left of $c$. In the latter case, however, $I_{f}$ includes the weak link, which by definition is worse than all edges between the culprits $c^{\prime}$ and $c$, including the edge $i-1 \rightarrow i$. This contradicts the observation preceding Lemma 13. A similar argument holds for $s_{j}$.

Let $C_{m}$ be the assignment on the set $S_{m}$ of middle strings (nodes) that has one cycle for each culprit, consisting of the greedy edges together with the back edge of the culprit. If we consider the application of the algorithm MGREEDY on the subset of strings $S_{m}$, it is easy to see that the algorithm will actually construct the assignment $C_{m}$. Theorem 10 then implies the following lemma:

Lemma 15. $\quad C_{m}$ is an optimal assignment on the set $S_{m}$ of middle strings.
Let the graph $G_{l}=\left(V_{l}, E_{l}\right)$ consist of the left/middle part of all blocks in the greedy path, that is, $V_{l}=S_{l} \cup S_{m}$ and $E_{l}$ is the set of non-weak greedy edges between nodes of $V_{i}$. Let $M_{i}$ be a maximum overlap assignment on $V_{l}$, as created by MGREEDY on the ordered sublist of edges in $V_{1} \times V_{1}$. Let $V_{r}=S_{m} \cup S_{r}$, and define similarly the graph $G_{r}=\left(V_{r}, E_{r}\right)$ and the optimal assignment $M_{r}$ on the right/middle strings. Let $l_{c}$ be the sum of the lengths of all culprit strings. Define $l_{1}=\sum_{i \in s_{i}} d\left(s_{i}, s_{i+1}\right)$ as the total length of all left extensions and $l_{r}=\sum_{i \in S_{t}} d\left(s_{i}^{R}, s_{i-1}^{R}\right)$ as the total length of all right extensions. (Here $x^{R}$ denotes the reversal of string $x$.) The length of the string produced by GREEDY is $l_{l}+l_{c}+l_{r}-o_{w}$, where $o_{w}$ is the summed block overlap (i.e., the sum of the overlaps of the weak links).

Denoting the overlap $\sum_{f \in E} o v(e)$ of a set of edges $E$ as $o v(E)$, define the cost of a set of edges $E$ on a set of strings (nodes) $V$ as

$$
\operatorname{cost}(E)=\|V\|-o v(E)
$$

Note that the distance plus overlap of a string $s$ to another equals $|s|$. Because an assignment (e.g., $M_{l}$ or $M_{r}$ ) has an edge from each node, its cost equals its distance weight. Since $V_{l}$ and $V_{r}$ are subsets of $S$ and $M_{l}$ and $M_{r}$ are optimal assignments, we have $\operatorname{cost}\left(M_{l}\right) \leq n$ and $\operatorname{cost}\left(M_{r}\right) \leq n$. For $E_{l}$ and $E_{r}$, we have that $\operatorname{cost}\left(E_{l}\right)=l_{l}+l_{c}$ and $\operatorname{cost}\left(E_{r}\right)=l_{r}+l_{c}$.

We have established the following (in)equalities:

$$
\begin{aligned}
l_{l}+l_{c}+l_{r}= & \left(l_{l}+l_{c}\right)+\left(l_{c}+l_{r}\right)-l_{c} \\
= & \operatorname{cost}\left(E_{l}\right)+\operatorname{cost}\left(E_{r}\right)-l_{c} \\
= & \left\|V_{l}\right\|-o v\left(E_{l}\right)+\left\|V_{r}\right\|-o v\left(E_{r}\right)-l_{c} \\
= & \operatorname{cost}\left(M_{l}\right)+o v\left(M_{l}\right)-o v\left(E_{l}\right)+\operatorname{cost}\left(M_{r}\right)+\operatorname{ov}\left(M_{r}\right) \\
& -\operatorname{ov}\left(E_{r}\right)-l_{c} \\
\leq & 2 n+\operatorname{ov}\left(M_{l}\right)-\operatorname{ov}\left(E_{l}\right)+o v\left(M_{r}\right)-o v\left(E_{r}\right)-l_{c} .
\end{aligned}
$$



FIG. 4. Left/middle and middle/right parts with weak links.

We proceed by bounding the overlap differences in the above equation. Our basic idea is to charge the overlap of each edge of $M$ to an edge of $E$ or a weak link or the back edge of a culprit in a way such that every edge of $E$ and every weak link is charged at most once and the back edge of each culprit is charged at most twice. This is achieved through combining the left/middle and middle/right parts carefully as shown below. For convenience, we will refer to the union operation for multisets (i.e., allowing duplicates) as the disjoint union.

Let $V$ be the disjoint union of $V_{l}$ and $V_{r}$, let $E$ be the disjoint union of $E_{l}$ and $E_{r}$, and let $G=(V, E)$ be the disjoint union of $G_{l}$ and $G_{r}$. Thus, each string in $S_{l} \cup S_{r}$ occurs once, while each string in $S_{m}$ occurs twice in $G$. We modify $E$ to take advantage of the block overlaps. Add each weak link to $E$ as an edge from the last node in the corresponding middle/right path of $G_{r}$ to the first node of the corresponding left/middle path of $G_{l}$. This procedure yields a new set of edges $E^{\prime}$. Its overlap equals $o v\left(E^{\prime}\right)=o v\left(E_{l}\right)+o v\left(E_{r}\right)+o_{w}$. A picture of $\left(V, E^{\prime}\right)$ for our previous example is given in Figure 4.

Let $M$ be the disjoint union of $M_{l}$ and $M_{r}$, an assignment on graph $G$. Its overlap equals $o v(M)=o v\left(M_{l}\right)+o v\left(M_{r}\right)$. Every edge of $M$ connects two $V_{l}$ nodes or two $V_{r}$ nodes; thus, all edges of $M$ satisfy the hypothesis of the following lemma.

Lemma 16. Let $N$ be any assignment on $V$. Let $e=t \rightarrow h$ be an edge of $N \backslash E^{\prime}$ that is not in $V_{r} \times V_{l}$. Then $e$ is dominated by either
(1) an adjacent $E^{\prime}$ edge, or
(2) a culprit's back edge with which it shares the head $h$ and $h \in V_{r}$, or
(3) a culprit's back edge with which it shares the tail $t$ and $t \in V_{l}$.

Proof. Suppose first that $e$ corresponds to a bad back edge. By Lemma 14, $h$ corresponds to a left node or to the first node of a culprit. In the latter case, $e$ is dominated by the back edge of the culprit (see the comment after Lemma 13). Therefore, either $h$ is the first node of a culprit in $V_{r}$ (and case (2) holds), or else $h \in V_{l}$. Similarly, either $t$ is the last node of a culprit in $V_{l}$ (and case (3) holds) or else $t \in V_{r}$. Since $e$ is not in $V_{r} \times V_{l}$, it follows then that case (2) or case (3) holds. (Note that, if $e$ is in fact the back edge of some culprit, then both cases (2) and (3) hold.)

Suppose that $e$ does not correspond to a bad back edge. Then $e$ must be dominated by some greedy edge since it was not chosen by GREEDY. If the greedy edge dominating $e$ is in $E^{\prime}$, then we have case (1). If it is not in $E^{\prime}$, then either $h$ is the first node of a culprit in $V_{r}$ or $t$ is the last node of a culprit in $V_{l}$, and in both cases $f$ is dominated by the back edge of the culprit. Thus, we have case (2) or (3).

Although Lemma 16 ensures that each edge of $M$ is bounded in overlap, it may be that some edges of $E^{\prime}$ are double charged. We modify $M$ without decreasing its overlap and without invalidating Lemma 16 into an assignment $M^{\prime}$ such that each edge of $E^{\prime}$ is dominated by one of its adjacent $M^{\prime}$ edges.

Lemma 17. Let $N$ be any assignment on $V$ such that $N \backslash E^{\prime}$ does not contain any edges in $V_{r} \times V_{l}$. Then there is an assignment $N^{\prime}$ on $V$ satisfying the following properties:
(1) $N^{\prime} \backslash E^{\prime}$ has also no edges in $V_{r} \times V_{l}$,
(2) $o u\left(N^{\prime}\right) \geq o v(N)$,
(3) each edge in $E^{\prime} \backslash N^{\prime}$ is dominated by one of its two adjacent $N^{\prime}$ edges.

Proof. Since $N$ already has the first two properties, it suffices to argue that if $N$ violates property (3), then we can construct another assignment $N^{\prime}$ that satisfies properties 1 and 2, and has more edges in common with $E^{\prime}$.

Let $e=k \rightarrow j$ be an edge in $E^{\prime}-N$ that dominates both adjacent $N$ edges, $f=i \rightarrow j$, and $g=k \rightarrow l$. By Lemma 7, replacing edges $f$ and $g$ of $N$ with $e$ and $i \rightarrow l$ produces an assignment $N^{\prime}$ with at least as large overlap. To see that the new edge $i \rightarrow l$ of $N^{\prime} \backslash E^{\prime}$ is not in $V_{r} \times V_{l}$, observe that if $i \in V_{r}$ then $j \in V_{r}$ because of the edge $f=i \rightarrow j\left(N \backslash E^{\prime}\right.$ does not have edges in $\left.V_{r} \times V_{t}\right)$, which implies that $k$ is in $V_{r}$ because of the $E^{\prime}$ edge $e=k \rightarrow j$ ( $E^{\prime}$ does not have edges in $V_{l} \times V_{r}$ ), which implies that also $l \in V_{r}$ because of the $N$ edge $g=k \rightarrow l$.

By Lemmas 16 and 17, we can construct from the assignment $M$ another assignment $M^{\prime}$ with at least as large total overlap, and such that we can charge the overlap of each edge of $M^{\prime}$ to an edge of $E^{\prime}$ or to the back edge of a culprit. Every edge of $E^{\prime}$ is charged for at most one edge of $M^{\prime}$, while the back edge of each culprit is charged for at most two edges of $M^{\prime}$ : for the $M^{\prime}$ edge entering the first culprit node in $V_{r}$ and the edge coming out of the last culprit node in $V_{l}$. Therefore, $o v(M) \leq o v\left(M^{\prime}\right) \leq o v\left(E^{\prime}\right)+2 o_{c}$, where $o_{c}$ is the summed overlap of all culprit back edges. Denote by $w_{c}$ the summed weight of all culprit cycles, that is, the weight of the (optimal) assignment $C_{m}$ on $S_{m}$ from Lemma 15. Then, $l_{c}=w_{c}+o_{c}$. As in the proof of Theorem 8, we have $o_{c}-2 w_{c} \leq n$ and $w_{c} \leq n$. (Note that the overlap of a culprit back edge is less than the length of the longest string in the culprit cycle.) Putting everything together, the string produced by GREEDY has length

$$
\begin{aligned}
l_{l}+l_{c}+l_{r}-o_{w} & \leq 2 n+o v\left(M_{l}\right)-o v\left(E_{l}\right)+o v\left(M_{r}\right)-o v\left(E_{r}\right)-l_{c}-o_{w} \\
& \leq 2 n+o v\left(M^{\prime}\right)-o v\left(E^{\prime}\right)-l_{c} \\
& \leq 2 n+2 o_{c}-l_{c} \\
& =2 n+o_{c}-w_{c} \\
& \leq 3 n+w_{c} \\
& \leq 4 n .
\end{aligned}
$$

## 6. Which Algorithm is the Best?

Having proved various bounds for the algorithms GREEDY, MGREEDY, and TGREEDY, one may wonder what this implies about their relative
performance. First of all, we note that MGREEDY can never do better than TGREEDY since the latter applies the GREEDY algorithm to an intermediate set of strings that the former merely concatenates.

Does the $3 n$ bound for TGREEDY then mean that it is the best of the three? This proves not always to be the case. In the example $\left\{c(a b)^{k},(a b)^{k+1} a,(b a)^{k} c\right\}$, GREEDY produces the shortest superstring $c(a b)^{k+1} a c$ of length $n=2 k+5$, whereas TGREEDY first separates the middle string to end up with something like $c(a b)^{k} a c(a b)^{k+1} a$ of length $4 k+6$.

Perhaps then GREEDY is always better than TGREEDY, despite the fact that we cannot prove as good an upper bound for it. This turns out not to be the case either, as shown by the following example. On input $\left\{c a b^{k}, a b^{k} a b^{k} a, b^{k} d a b^{k-1}\right\}$, TGREEDY separates the middle string, merges the other two, and next combines these to produce the shortest superstring $c a b^{k} d a b^{k} a b^{k} a$ of length $3 k+6$, whereas GREEDY merges the first two, leaving nothing better than $c a b^{k} a b^{k} a b^{k} d a b^{k-1}$ of length $4 k+5$.

Another greedy type of algorithm that may come to mind is one that arbitrarily picks any of the strings and then repeatedly merges on the right the string with maximum overlap. This algorithm, call it NAIVE, turns out to be disastrous on examples like
$\{a b c d e, b c d e \# a, c d e \# a \# b, d e \# a \# b \# c, e \# a \# b \# c \# d, \# a \# b \# c \# d \# e\}$.
Instead of producing the optimal abcde\#a\#b\#c\#d\#e, NAIVE might pick $\# a \# b \# c \# d \# e$ as a starting point to produce
$\# a \# b \# c \# d \# e \# a \# b \# c \# d e \# a \# b \# c d e \# a \# b c d e \# a b c d e$.
It is clear that in this way superstrings may be produced whose length grows quadratically in the optimum length $n$.

## 7. Lower Bound

We show here that the superstring problem is MAX SNP-hard. This implies that if there is a polynomial time approximation scheme for the superstring problem, then there is one also for a wide class of optimization problems, including several variants of maximum satisfiability, the node cover and independent set problems in bounded-degree graphs, max cut, etc. This is considered rather unlikely. ${ }^{1}$

Let $A, B$ be two optimization (maximization or minimization) problems. We say that $A$ L-reduces (for linearly reduces) to $B$ if there are two polynomial time algorithms $f$ and $g$ and constants $\alpha$ and $\beta>0$ such that:
(1) Given an instance $a$ of $A$, algorithm $f$ produces an instance $b$ of $B$ such that the cost of the optimum solution of $b, \operatorname{opt}(b)$, is at most $\alpha \cdot \operatorname{opt}(a)$, and
(2) Given any solution $y$ of $b$, algorithm $g$ produces in polynomial time a solution $x$ of $a$ such that $|\operatorname{cost}(x)-\operatorname{opt}(a)| \leq \beta|\operatorname{cost}(y)-\operatorname{opt}(b)|$.
Some basic facts about L-reductions are: First, the composition of two L-reductions is also an L-reductions. Second, if problem $\mathcal{A}$ L-reduces to problem $B$ and $B$ can be approximated in polynomial time with relative error $\epsilon$

[^1](i.e., within a factor of $1+\epsilon$ or $1-\epsilon$ depending on whether $B$ is a minimization or maximization problem), then $\boldsymbol{A}$ can be approximated with relative error $\alpha \beta \epsilon$. In particular, if $B$ has a polynomial time approximation scheme, then so does $A$. The class MAXSNP is a class of optimization problems defined syntactically in Papadimitriou and Yannakakis [1988]. It is known that every problem in this class can be approximated within some constant factor. A problem is MAXSNP-hard if every problem in MAXSNP can be L-reduced to it.

## Theorem 18. The superstring problem is MAX SNP-hard.

Proof. The reduction is from a special case of the TSP with triangle inequality. Let TSP $(1,2)$ be the TSP restricted to instances where all the distances are either 1 or 2 . We can consider an instance to this problem as being specified by a graph $H$; the edges of $H$ are precisely those that have length 1 while the edges that are not in $H$ have length 2 . We need here the version of the TSP where we seek the shortest Hamiltonian path (instead of cycle), and, more importantly, we need the additional restriction that the graph $H$ be of bounded degree (the precise bound is not important). It was shown in Papadimitriou and Yannakakis [1993] that the $\operatorname{TSP}(1,2)$ problem (even for this restricted version) is MAX SNP-hard.

Let $H$ be a graph of bounded degree $D$ specifying an instance of $\operatorname{TSP}(1,2)$. The hardness result holds for both the symmetric and the asymmetric TSP (i.e., for both undirected and directed graphs $H$ ). We let $H$ be a directed graph here. Without loss of generality, assume that each vertex of $H$ has outdegree at least 2. The reduction is similar to the one of Gallant et al. [1980] used to show the NP-completeness of the superstring decision problem. We have to prove here that it is an L-reduction. For every vertex of $v$ of $H$, we have two letters $v$ and $v^{\prime}$. In addition, there is one more letter \#. Corresponding to each vertex $v$ we have a string $v \# v^{\prime}$, called the connector for $v$. For each vertex $v$, enumerate the edges out of $v$ in an arbitrary cyclic as $\left(v, w_{0}\right), \ldots,\left(v, w_{d-1}\right)\left({ }^{*}\right)$. Corresponding to the $i$ th edge $\left(v, w_{i}\right.$ ) out of $v$, we have a string $p_{i}(v)=$ $v^{\prime} w_{i-1} v^{\prime} w_{i}$, where subscript arithmetic is modulo $d$. We say that these strings are associated with $u$.

Let $n$ be the number of vertices and $m$ the number of edges of $H$. If all vertices have degree at most $D$, then $m \leq D n$. Let $k$ be the minimum number of edges whose addition to $H$ suffices to form a Hamiltonian path. Thus, the optimal cost of the TSP instance is $n-1+k$. We shall argue that the length of the shortest common superstring is $2 m+3 n+k+1$. It will follow then that the reduction is linear since $m$ is linear in $n$.

Consider the distance-weighted graph $G_{S}$ for this set of strings, and let $G_{2}$ be its subgraph with only edges of minimal weight (2). Clearly, $G_{2}$ has exactly one component for each vertex of $H$, which consists of a cycle of the associated $p$ strings, and a connector that has an edge to each of them. We need only consider "standard" superstrings in which all strings associated with some vertex form a subgraph of $G_{2}$, so that only the last $p$ string has an outgoing edge of weight more than 2 ( 3 or 4). Namely, if some vertex fails this requirement, then at least two of its associated strings have outgoing edges of weight more than 2 , thus we do not increase the length by putting all its $p$ strings directly after its connector in a standard way. A standard superstring naturally corresponds to an ordering of vertices $v_{1}, v_{2}, \ldots, v_{n}$.

For the converse, there remains a choice of which string $q$ succeeds a connector $v_{i} \# v_{i}^{\prime}$. If $H$ has an edge from $v_{i}$ to $v_{i+1}$ and the "next" edge out of $v_{i}$ (in $\left(^{*}\right)$ ) goes to, say $v_{j}$, then choosing $q=v_{i}^{\prime} v_{i+1} v_{i}^{\prime} v_{j}$ results in a weight of 3 on the edge from the last $p$ string to the next connector $v_{i+1} \# v_{i+1}^{\prime}$, whereas this weight would otherwise be 4 . If $H$ doesn't have this edge, then the choice of $q$ doesn't matter. Let us call a superstring "Standard" if in addition to being standard, it also satisfies this latter requirement for all vertices.

Now suppose that the addition of $k$ edges to $H$ gives a Hamiltonian path $v_{1}, v_{2}, \ldots, v_{n-1}, v_{n}$. Then, we can construct a corresponding Standard superstring. If the out-degree of $v_{i}$ is $d_{i}$, then its length will be $\sum_{i=1}^{n}\left(2+2 d_{i}+1\right)+$ $k+1=3 n+2 m+k+1$.

Conversely, suppose we are given a common superstring of length $3 n+2 m$ $+k+1$. This can then be turned into a Standard superstring that is no longer. If $v_{1}, v_{2}, \ldots, v_{n}$ is the corresponding order of vertices, it follows that $H$ cannot be missing more than $k$ of the edges ( $v_{i}, v_{i+1}$ ).

Since the strings in the above L-reduction have bounded length (4), the reduction applies also to the maximization version of the superstring problem [Tarhio and Ukkonen, 1988; Turner, 1989]. That is, maximizing the total compression is also MAX SNP-hard.

## 8. Open Problems

We end the paper with several open questions raised from this research:
(1) Obtain an algorithm that achieves a performance better than 3 times the optimum.
(2) Prove or disprove the conjecture that GREEDY achieves 2 times the optimum.
acknowledgments. We thank Samir Khuller and Vijay Vazirani for discussions on the superstring algorithms (Samir brought the authors together), and Rafi Hassin for bringing Hoffman's and others' work on Monge sequences to our attention. We would also like to thank the referees for their helpful comments.

## REFERENCES

alon, N., Cosares, S., Hochbaum, D., and Shamir, R. 1989. An algorithm for the detection and construction of Monge Sequences. Lin. Alg. Appl. 114/115, 669-680.
Arora, A., Lund, C., Motwani, R., Sudan, M., and Szegedy, M. 1992. Proof verification and hardness of approximation problems. In Proceedings of the 33rd IEEE Symposium on the Foundations of Computer Science. IEEE, New York, pp. 14-23.
Barnes, E., and Hoffman, A., 1985. On transportation problems with upper bounds on leading rectangles. SLAM J. Alg. Disc. Meth. 6, 487-496.
Fine, N., and Wilf, H. 1965. Uniqueness theorems for periodic functions. Proc. Amer. Math. Soc. 16, 109-114.
Gallant, J., Maier, D., and Storer, J. 1980. On finding minimal length superstrings. $J$. Comput. Syst. Sci. 20, 50-58.
Garey, M., and Johnson, D. 1979. Computers and Intractability. Freeman, New York.
Hoffman, A., 1963. On simple transportation problems. In Convexity: Proceedings of Symposia in Pure Mathematics, vol. 7. American Mathematical Society, Providence, R.I., pp. 317-327.
Lesk, A. (ED) 1988. Computational Molecular Biology, Sources and Methods for Sequence Analysis. Oxford University Press.

## Linear Approximation of Shortest Superstrings

LI, M. 1990. Towards a DNA sequencing theory. In Proceedings of the 31st IEEE Symposium on Foundations of Computer Science. IEEE, New York, pp. 125-134.
Papadimitriou, C., and Steiglitz, K. 1982. Combinatorial Optimization: Algorithms and Complexity. Prentice-Hall, Englewood Cliffs, N.J.
Papadimitriou, C., and Yannakakis, M. 1988. Optimization, approximation, and complexity classes. In Proceedings of the 20th ACM Symposium on Theory of Computing (Chicago, Ill., May 2-4). ACM, New York, pp. 229-234.
Papadimitriou, C. and Yannakakis, M. 1993. The traveling salesman problem with distances one and two. Math. Oper. Res. 18, 1, 1-11.
Peltola, H., Soderlund, H., Tarhio, J., and Ukkonen, E. 1983. Algorithms for some string matching problems arising in molecular genetics. In Information Processing 83 (Proceedings of IFIP Congress, 1983). Elsevier Science Publishers R. V. (North-Holland), Amsterdam, The Netherlands, pp. 53-64.
Storer, J. 1988. Data compression: methods and theory. Computer Science Press, Rockville, Md.
TARHIO, J., AND UkKONEN, E. 1988. A Greedy approximation algorithm for constructing shortest common superstrings. Theoret. Comput. Sci. 57, 131-145.
TURNER, J., 1989. Approximation algorithms for the shortest common superstring problem. Inf. Comput. 83, 1-20.
Valiant, L. G. 1984. A Theory of the learnable. Commun. ACM 27, 11 (Nov.), 1134-1142.
RECEIVED JULY 1991; REVISED DECEMBER 1992; ACCEPTED JANUARY 1993


[^0]:    A. Blum was supported by a National Science Foundation (NSF) Graduate Fellowship. Part of this work was done while A. Blum was visiting AT\&T Bell Labs.
    T. Jiang was supported in part by a grant from SERB, McMaster University and NSERC Operating Grant OGP0046613.
    M. Li was supported in part by the NSERC Operating Grants OGP0036747 and OGP0046506.
    J. Tromp was supported in part by NSERC Grant OGP0036747 while the author was visiting at the University of Waterloo.
    Authors' present addresses: A. Blum, School of Computer Science, Carnegie-Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213. E-mail: avrim@theory.cs.cmu.edu; T. Jiang, Department of Computer Science, McMaster University, Hamilton, Ont., Canada L8S 4K1. E-mail: jiang (amaccs. momaster.ca; M. Li, Department of Computer Science, University of Waterloo, Waterloo, Ont., Canada N3L 3G1. E-mail: mli@watmath.waterloo.edu; J. Tromp, CWI, P.O. Box 4079, 1009 AB Amsterdam, The Netherlands. E-mail: tromp@cwi.nl; M. Yannakakis, Room 2C-319, AT\&T Bell Labs, 600 Mountain Avenue, Murray Hill, NJ 07974. E-mail: mihalis@research.att.com.
    Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

[^1]:    ${ }^{1}$ In fact, Arora et al. [1992] have recently shown that MAX SNP-hard problems do not have polynomial time approximation schemes, unless $P=N P$.

