Exact Algorithms for the Master Ring Problem*

Hadas Shachnai[†] Lisa Zhang [‡] Tomomi Matsui [§]

October 11, 2006

Abstract

We consider the *master ring problem (MRP)* which often arises in optical network design. Given a network which consists of a collection of interconnected rings $R_1, \ldots, R_K$, with $n_1, \ldots, n_K$ distinct nodes, respectively, we need to find an ordering of the nodes in the network that respects the ordering of every individual ring, if one exists. We show that MRP is NP-complete, therefore, it is unlikely to be solvable by a polynomial time algorithm. Our main result is an algorithm which solves MRP in $Q \cdot \prod_{k=1}^{K}(n_k/\sqrt{2})$ steps, for some polynomial Q, as the n_k values become large. For the *ring clearance problem*, a special case of practical interest, our algorithm achieves this running time for rings of *any* size $n_k \geq 2$. This yields the first nontrivial improvement, by factor of $(2\sqrt{2})^K \approx (2.82)^K$, over the running time of the naive algorithm, which exhaustively enumerates all $\prod_{k=1}^{K}(2n_k)$ possible solutions.

Keywords: Master ring, shortest common supersequence, optical networks, exact algorithms.

1 Introduction

1.1 Problem Statement and Motivation

A ring is a popular network topology due to its simplicity and effectiveness in fault recovery. To carry a demand between two nodes on a ring, traffic can be routed simultaneously clockwise and counter-clockwise, one as the primary path and the other as the backup path. The SONET (*Synchronous Optical NETwork*) technology further popularized rings [13]. Often an optical network consists of a collection of interconnecting rings. A *master ring* contains every node in the network exactly once and respects the node ordering of every individual ring. The *master ring problem (MRP)* is to find such a ring, whenever it exists.

Formally, the master ring problem is defined as follows. Suppose that a network consists of K rings, $R_1, \ldots, R_K$, with $n_1, \ldots, n_K$ distinct nodes, respectively. Each ring has two *orientations*, clockwise and counter-clockwise. We say that R is a *subring* of M (or M is a *master ring* of R) if

*A preliminary version of this paper appears in [14].

†Computer Science Dept., Technion, Haifa 32000, Israel.

‡Bell Labs, Lucent Technologies, 600 Mountain Ave., Murray Hill, NJ 07974.

§Dept. of Math. Info., Grad. School of Info. Sci. and Tech., The Univ. of Tokyo, Tokyo 113-8656, Japan.

either the clockwise or the counter-clockwise orientation of R can be obtained from M by erasing zero or more nodes from M which contains distinct nodes. We assume if two rings R_i and R_j intersect then they have at least two nodes in common. This is because two common nodes can tolerate one node failure when supporting a demand between a node in R_i and a node in R_j. Our goal is to find a master ring whenever it exists.

Consider an instance of MRP as shown in Figure 1. The network consists of 3 rings. R_1 has the nodes $abcdef$, R_2 has the nodes $achg$, and R_3 has the nodes $ghcdi$. A possible master ring is $abghcdefi$.

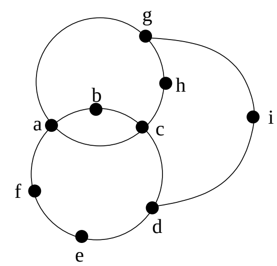

Figure 1: An instance of MRP.

There are a number of reasons for finding master rings. For example, as a network evolves with growing traffic, it expands from an initially small number of rings, to include a large collection of rings. Unfortunately, such expansion is often carried out in an ad-hoc manner, with circuits added and torn down over time. As a result, the network may have unnecessarily complex topology that makes network management a nightmare. To replace a spaghetti-like network, one simple topology is a master ring. Since a master ring respects the node ordering of every existing ring, it has the advantage of preserving the routing label of every demand intra to an existing ring. Indeed, a demand may traverse more nodes around the master ring than around its original ring; however, preserving the order in which the nodes are traversed causes less interruption. Even if the network is not sought to be rebuilt, it still needs to handle the routine downtime, for purposes such as software upgrade. A master ring can then serve as a simple backup topology. Providing a master ring (whenever possible) to a network management system simplifies its operation and is therefore valuable [12, 1, 2]. We are not aware of any earlier analysis on MRP. As far as we know, in practice MRP is solved via brute-force enumeration.

We emphasize that the master ring can be viewed as a “logic” ring. That is, two neighboring nodes in the ring do not need to be physically connected by links already existing in the network. (Such links can be added once the master ring is set as a new/backup topology.) We also remark that, for long-term network optimization tasks such as network rebuilding, we often have the luxury of finding an exact solution at the cost of longer running time.

1.2 Sequence Representation

One convenient way to represent the rings is to use sequences. Each orientation of a ring with n nodes corresponds to n *sequences*, depending on the node with which the sequence starts. Figure 2 shows the sequence representation of the instance in Figure 1. For example, the ring R_1 in Figure 2 has 6 clockwise sequences: $abcdef$, $bcdefa$, $cdefab$, $defabc$, $efabcd$, $fabcde$ and 6 counter-clockwise sequences: $fedcba$, $edcbaf$, $dcbafe$, $cbafed$, $bafedc$, $afedcb$. We also refer to each sequence as an *opening* of a ring. We say that S is a *subsequence* of T (or T is a *supersequence* of S) if S can be obtained from T by erasing zero or more symbols from T. Therefore, R is a *subring* of M (or M is a *master ring* of R) if some sequence that corresponds to R is a subsequence of a sequence that corresponds to M (see Figure 2).

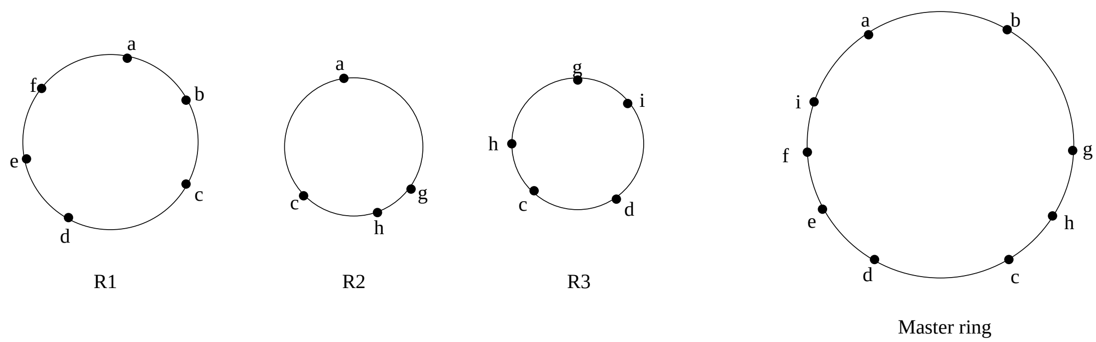

Figure 2: (Left) Three rings R_1, R_2 and R_3. (Right) A possible master ring. For example, R_1 is a subring since its clockwise sequence $abcdef$ is a subsequence of the sequence $abghcdefi$ corresponding to the master ring; R_2 is a subring since its clockwise sequence $aghc$ is a subsequence; R_3's counter-clockwise sequence $ghcdi$ is a subsequence.

Given K sequences, each with any symbol appearing at most once, we note that it is easy to find a supersequence that contains each symbol once, if one exists. We construct a directed graph $G = (V, E)$, whose vertex set consists of the symbols in the K sequences and whose edge set consists of directed edges of the form (a, b), where a appears immediately before b in a sequence. Recall that a topological sort of an acyclic digraph is any linear order on the vertices respecting the graph's partial order. Hence, if G is acyclic then a topological sort of G is a (minimum possible length) supersequence of all K sequences. Deciding whether a digraph is acyclic and finding a topological sort are polynomially solvable (see e.g. [4]). Thus, our main task is to determine a set of sequences which can be represented as an acyclic digraph (whenever such a set exists). This is the focus of the paper.

1.3 Results

In Section 5 we show that MRP is NP-complete, therefore, it is unlikely to be solvable by a polynomial time algorithm. Our main result (in Section 2) is an algorithm for MRP whose running time approaches $Q \cdot \prod_{k=1}^{K}(n_k/\sqrt{2})$ for some Q that is polynomial in the input size, as the n_k values become large. For the *ring clearance problem*, a special case of practical interest where one ring intersects with all the other rings, our algorithm achieves this running time for rings of *any* size $n_k \geq 2$ (see in Section 3). This yields the first non-trivial improvement, by factor of $(2\sqrt{2})^K \approx 2.82^K$, over the naive algorithm which exhaustively enumerates all $\prod_{k=1}^{K}(2n_k)$ possible solutions (see, e.g., in [2]).

Our algorithm applies enumeration guided by an *intersection graph* of the network, which represents the interconnections among the rings. The graph is used for identifying subsets of rings whose openings leave only a few consistent openings for all other rings, thereby decreasing the remaining number of enumeration steps. While enumeration alone is inefficient, and using the intersection graph alone may result in a false solution for our problem (see in Section 2), we show that combining the two yields a significant improvement in running time, and guarantees that a master ring will be found, if one exists. We believe that similar techniques can be used in solving exactly other related problems, such as shortest common supersequence (SCS) and feedback arc set (FAS) and their variants (see in Section 4).

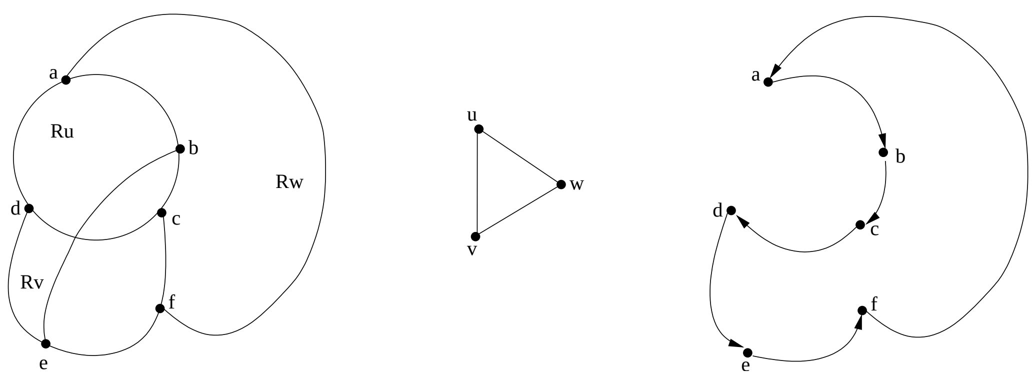

Figure 3: (Left) An instance of MRP: R_u consists of nodes $abcd$, R_v consists of $cdef$ and R_w consists of $befa$. (Middle) The intersection graph H. (Right) R_u, R_v and R_w induce a large ring.

2 Algorithm

A naive solution for MRP is to enumerate all possible sequences for each ring and find if there is a topological sort for each resulting directed graph. Obviously, trying the total of $\prod_{1\leq k\leq K}(2n_k)$ possibilities suffices for solving the problem; the running time is $P \cdot \prod_{1\leq k\leq K}(2n_k)$, where P is the polynomial time required for topological sort. We describe below an algorithm which avoids enumerating some of these possibilities, by using the *intersection graph* of the network.

Before we apply our algorithm, we first eliminate all *singleton* nodes from each ring, i.e., those nodes that appear only in one ring. If node a is a singleton, then a can be ignored when constructing the master ring. Indeed, if a master ring exists without a, then a may always be added to the master ring. From now on we may assume without loss of generality that every node appears in at least 2 rings.

We construct an undirected intersection graph H that shows how the rings are interconnected. The graph H consists of K vertices, each corresponding to one of the K rings. If two rings share common nodes, then there is an edge between their corresponding vertices in H. For clarity, we use *vertices* and *edges* when referring to the elements in the graph H and *nodes* and *links* – when referring to the elements of a ring. We also use letters near the beginning of the alphabet (such as a, b, c and d) when referring to nodes in a ring and letters near the end of the alphabet (such as u, v and w) when referring to vertices in H. For a vertex u in H, we use R_u to represent the corresponding ring. (See Figure 3 for an example.) Our algorithm is motivated by the following

Algorithm $\mathcal{A}_{MR}$

0 Eliminate singleton nodes from $R_1, \ldots, R_K$. Construct the graph H with vertex set V.

Phase 1. Low-degree vertices

1 $N = L = \emptyset$; $\delta = \Theta(\log n)$

2 While there is $v \in V - L - N$, where v has degree lower than δ
  add vertex v to set L and its neighbors to set N.

3 For $u \in N$, try all possible sequences for R_u.

4 For $v \in L$ with x neighbors, try at most $2x$ possible sequences for R_v.
(See Lemma 1.)

Phase 2. Dominating set

5 Find a dominating vertex set D for the vertices $v \in H$ such that $v \in V - L - N$.

6 For $u \in D$, try all possible sequences for R_u.

Phase 3. Remaining vertices

7 Let $C = V - L - N - D$.

8 For $u \in C$, try a total of $y^{|C|}$ combinations of sequences for R_u,
where y is given in Lemma 5.

9 For each combination of sequences for vertices in $N \cup L \cup D \cup C$,
  find a supersequence T using topological sort.

10 If T exists, a master ring is found. Algorithm terminates.

11 Output no master ring exists.

Figure 4: The master ring algorithm $\mathcal{A}_{MR}$.

observation. Consider a vertex u in H. If R_v is already opened (i.e., our algorithm has found an opening for R_v), and v is a neighbor of u, then the number of *consistent* openings of R_u is limited. We say that a set of sequences is *consistent* if they have a supersequence. For example, suppose that R_u and R_v have in common the nodes a and b, and R_v orders a before b; then, R_u would have to as well.

We note, however, that even if any two *neighboring* rings have consistent openings, it does not necessarily imply consistent openings for all rings. Consider the instance of Figure 3. When R_u is oriented clockwise, and R_v, R_w are oriented counter-clockwise, they induce a large ring $abcdef$. If R_u has opening $abcd$, R_v has opening $cdef$, and R_w has opening $efab$ then no opening of this induced ring contains the three sequences as subsequences. Therefore, these three openings cannot be consistent with one another. However, any two of these openings are consistent. (See Figure 3, Right). If, instead, R_w has the opening $abef$, then the three openings are consistent and have a master ring $abcdef$. The example in Figure 3 shows that we cannot use the graph H alone for determining good openings for all rings, since this graph indicates only the 'local' dependencies among the rings. To guarantee that no induced rings remain in the network after we open $R_1, \ldots, R_K$, we use the properties of the graph H only as guidance for the algorithm.

In our algorithm, $\mathcal{A}_{MR}$, we identify a low-degree vertex u in H and enumerate all possible openings of u's neighbors. Since u has low degree, relatively few rings are opened, but this dramatically limits the number of consistent openings of R_u (see Lemma 1). Recall that a *dominating set* for an undirected graph $G = (V, E)$ is a subset of vertices $V' \subseteq V$ such that any vertex $v \in V \setminus V'$ is a neighbor of some vertex in V'. When H has only high-degree vertices, we can find in the graph a small dominating set. By enumerating all possible openings for the (small number of) vertices in the dominating set, we can reduce the number of consistent openings for each remaining vertex by a constant factor (see Lemma 5). In our algorithm, we define low and high degree vertices through a parameter δ (see below); If a vertex $u \in H$ has degree lower than δ then u is a *low-degree* vertex. A pseudocode of algorithm $\mathcal{A}_{MR}$ is given in Figure 4.

3 Analysis

For simplicity of exposition, we assume throughout the analysis that all of the rings are of the same size, n, and set $\delta = \log n/c$, where $c \geq 3$ is some constant. Later, we show how the analysis extends to rings of arbitrary sizes.

In the following we show the correctness and worst-case running time of algorithm $\mathcal{A}_{MR}$. For correctness, if the algorithm finds a sequence T that is a supersequence for some opening of every ring R_k, where $1 \leq k \leq K$, the master ring can be defined by T. However, since our algorithm does not exhaustively enumerate all of the $2n$ openings of each ring, if it does not find a supersequence we need to verify that we have not missed any opening that could have lead to a supersequence. We start by analyzing the first phase.

Lemma 1 *If a vertex $v \in L$ has x neighbors in H, and each neighbor is opened, then at most $2x$ sequences of v can be consistent with the x neighboring sequences.*

Proof: Let u be a neighbor of v, and let S_u be the sequence representing the opening of the ring R_u. Consider the subsequence T_u of S_u that consists of the nodes common to R_u and R_v. Let a_u be the first symbol in T_u. Since we have no singleton nodes, we know that $\bigcup_u T_u$ contains all the nodes in R_v. Therefore, if S_v begins with a node in T_u for some neighbor u, then S_v has to begin with a_u; otherwise, S_v cannot be consistent with S_u since T_u contains at least two nodes. There are two possible directions starting at a_u. □

From Lemma 1, in Line 4 of the algorithm we try at most 2δ sequences for any ring R_v, such that v is a low-degree vertex in H. This allows us to bound the running time of Phase 1.

Lemma 2 *The running time of Phase 1 is at most $(1 + o(1))\frac{n^{\hat{k}}}{2^{\hat{k}(c-2)}}$ where $\hat{k} = |L| + |N|$ is the total number of vertices dealt with in this phase.*

Proof: It is easy to see that the number of combinations that Phase 1 tries is bounded by $(2n)^{|N|}(2\delta)^{|L|}$. However, to execute Line 4, we need to determine the $x < \delta$ potential nodes to open a ring. This can be done in time $O(\alpha\delta|L|)$ for some constant α, using the procedure described in Lemma 1. Therefore, the outer loop in our algorithm, Phase 1, takes at most $(2n)^{|N|}(\alpha\delta|L|+(2\delta)^{|L|})$, which is $(1 + o(1))(2n)^{|N|}(2\delta)^{|L|}$. Note that the $o(1)$ term is a function of the ring size n and c.

Let us look at the running time more closely. Let $\hat{k} = |L| + |N|$ be the total number of vertices handled in Phase 1. Since L consists of vertices with degree lower than δ, we have $|L| \geq \hat{k}/\delta$ and $|N| \leq \hat{k}(1 - 1/\delta)$. Therefore,

$$\begin{aligned}
(2n)^{|N|}(2\delta)^{|L|} &= n^{|L|+|N|}\left(\frac{2\delta}{n}\right)^{|L|} 2^{|N|} \\
&\leq n^{\hat{k}}\left(\frac{2\delta}{n}\right)^{\hat{k}/\delta} 2^{\hat{k}(1-\frac{1}{\delta})} \\
&= n^{\hat{k}} 2^{\hat{k}(\frac{\log 2\delta}{\delta}-c)} 2^{\hat{k}(1-\frac{1}{\delta})} \\
&\leq \frac{n^{\hat{k}}}{2^{\hat{k}(c-2)}}.
\end{aligned}$$

□

Let us bound the size of the dominating set D in Phase 2.

Lemma 3 $|D| \leq |V - L - N| \cdot \frac{1+\ln(1+\delta)}{1+\delta}$.

Proof: We first prove the next claim, which generalizes a result in [3]. Let $G = (V, E)$ be a graph, such that all the vertices in $V' \subseteq V$ have degree at least s. Then there exists a subset of vertices $V'' \subseteq V'$ of size at most $|V'|\frac{1+\ln(1+s)}{1+s}$, such that $U = (V \setminus V') \bigcup V''$ is a dominating set for G.

Consider the following Greedy algorithm. (i) We start by adding all the vertices in $V \setminus V'$ to U. (ii) Let W be the set of vertices in V' that are not in U and do not have a neighbor in U; While

$|W| > |V'|/(s+1)$ do: Find a vertex $v \in W$ such that v has a maximal number of neighbors in W; add v to U. (*iii*) Add the vertices in W to U.

Clearly, U is a dominating set for G. To bound the size of V'', we first note that, by an averaging argument, since all the vertices in $V \setminus V'$ are added to U, the number of iterations until $|W| \leq |V'|/(s+1)$ is at most $|V'| \ln(s+1)/(s+1)$. (A similar argument is given in the analysis of the deterministic algorithm for the dominating set problem in [3]; we omit the details.) Hence, we get that $|V''| \leq |V'| \ln(s+1)/(s+1) + |V'|/(s+1)$. If we set $V' = V - N - L$ and $s = \delta$, then $U = N \cup L \cup V''$ is a dominating set for H; thus, V'' is a dominating set for $H \setminus (N \cup L)$, and the lemma follows. □

The greedy algorithm in Lemma 3 takes time at most quadratic in K, the number of vertices in H. Hence,

Lemma 4 *The running time of Phase 2 is* $poly(K) + (2n)^{|D|}$.

We now discuss how to efficiently find openings for the remaining vertices in C during Phase 3.

Lemma 5 *The running time of Phase 3 is at most* $poly(K)y^{|C|}$*, where* $y \leq \sqrt{2}(2 + n/2)$*. Hence, the running time of phase 3 is* $poly(K)(n/\sqrt{2})^{|C|}$.

Proof: During Phase 3, every vertex $u \in C$ has some neighbor v in the dominating set D. By assumption, R_u and R_v have at least 2 nodes, say a and b, in common. Any sequence of R_v defines an ordering of a and b, i.e., a appears before b or after b. Among the $2n$ sequences of R_u, exactly n respect this ordering of a and b. Any of the other n sequences that disrespect the ordering cannot produce a topological sort and therefore need not be considered. We get that it suffices to enumerate at most n sequences for the ring R_u, for any $u \in C$.

We can further reduce the number of enumerations using the *pairing algorithm* described below. Instead of directly enumerating n possible sequences for R_u where $u \in C$, we pair up the sequences so that one sequence in a pair begins with a node, say a, and the other sequence in the pair ends with the node a. We refer to a as the *pivot* of the pair.

We first argue at most 2 out of the n sequences cannot be paired up with another sequence. In the clockwise direction, let $\alpha a \beta b$ be the sequence that has a before b and finishes at b. Here α is the subsquence before a and β is the subsequence between a and b. Note that α or β could be empty. It is easy to see that there are $|\alpha| + 1$ shifts in the clockwise direction that have a before b and $\lfloor \frac{|\alpha|+1}{2} \rfloor$ pairs can be paired. In the counter-clockwise direction, $\bar{\beta} a \bar{\alpha} b$ is the sequence that has a before b and finishes at b. Here $\bar{\alpha}$ and $\bar{\beta}$ are reverses of α and β. There are $|\beta| + 1$ shifts in the counter-clockwise direction that have a before b and $\lfloor \frac{|\beta|+1}{2} \rfloor$ pairs can be paired. In total, at least $|\alpha| + |\beta|$ sequences are paired, which leaves 2 unpaired.

More concretely, let us consider the following example. For $u \in C$, suppose the ring R_u has 6 nodes $abcdef$ clockwise. Suppose that u's neighbor $v \in D$ has chosen a sequence for R_v in which b precedes e. Therefore, any sequence for R_u needs to have b before e, else there is no topological sort. Among the 12 possible sequences for R_u, the following 6 have b before e.

$fabcde$ $abcdef$ $bcdefa$ $dcbafe$ $cbafed$ $bafedc$

We pair up the 1st and 2nd sequences, $fabcde$ and $abcdef$, with pivot f, and pair up the 4th and 5th sequences, $dcbafe$ and $cbafed$ with pivot d. The 3rd sequence, $bcdefa$, and the 6th sequence, $bafedc$, remain singletons.

We first enumerate the $2 + n/2$ groups (at most $n/2$ pairs and at most 2 unpaired singletons) for every $u \in C$. This gives a total of at most $(2 + n/2)^{|C|}$ possibilities. In the following we show that to determine the actual sequence within each pair for any $u \in C$, we do not need to try both possibilities. In fact, a total of $2^{|C|/2}$ trials suffices. Hence, the total number of trials is $(2 + n/2)^{|C|}2^{|C|/2}$, which implies $y \leq \sqrt{2}(2 + n/2)$.

For example, suppose ring R_u, where $u \in C$, has the above 6 possibilities, and we are considering the first pair with pivot a. Since a is not a singleton node, a necessarily appears in another ring, say R_w. If the sequence for R_w is decided, then we necessarily know which sequence, $abcdef$ or $bcdefa$, would be good. This is because R_u and R_w must share a node other than a; let us call this node c. Since a is a pivot for R_u, if a appears before c for R_w, then only the first sequence $abcdef$ can be good; if a appears after c for R_w, then only the second sequence $bcdefa$ can be good.

In general, we construct a directed graph, F, where each vertex corresponds to a vertex in C. We put a directed edge from u to w if the pivot of u is a vertex in the ring R_w. If there are multiple such rings R_w for u we choose an arbitrary one. As argued above, if there is a directed edge from u to w, then we only need to enumerate the two choices in a chosen pair for R_w, and the choice for R_u is implied. We determine which rings to enumerate as follows. We mark a *cross* on a vertex to indicate that the choice is implied and we mark a *circle* on a vertex to indicate that we enumerate both possibilities. Initially, we mark a cross on a vertex u if it has no outgoing edges. This means the pivot of u appears in some ring R_w that belongs to $L \cup N \cup D$. Hence, the sequence for R_w is already chosen, and therefore the sequence for R_u is implied. For each vertex u in F that is not yet marked, we follow the directed edges, starting from u, until (i) we have reached a marked vertex (either with a circle or with a cross), or (ii) we stop right before the path from u intersects itself, i.e., in a vertex z such that there is an edge (z, u). In the latter case, we circle the vertex where we stop. In both cases, we also mark a cross on every (unmarked) vertex along the path. (See Figure 5.)

It is easy to verify that the choice for each vertex with a cross can be implied from the choice for some vertex with a circle. In terms of the running time, we observe that at most half of the vertices in F can be circled, since each circled vertex needs at least one distinct vertex that has a cross. To mark each vertex in F with a circle or cross requires visiting each vertex once. Hence, the time requirement is linear in $|C|$. It follows that the running time of Phase 3 is at most $poly(|C|)(2 + n/2)^{|C|} \cdot 2^{|C|/2}$, which is $poly(K)(n/\sqrt{2}))^{|C|}$. □

From Lemmas 1 and 5 we see that although algorithm $\mathcal{A}_{MR}$ does not enumerate all possibilities for the vertices in L and C we do not miss out any potentially good opening. Our algorithm is therefore correct. We bound the running time as follows.

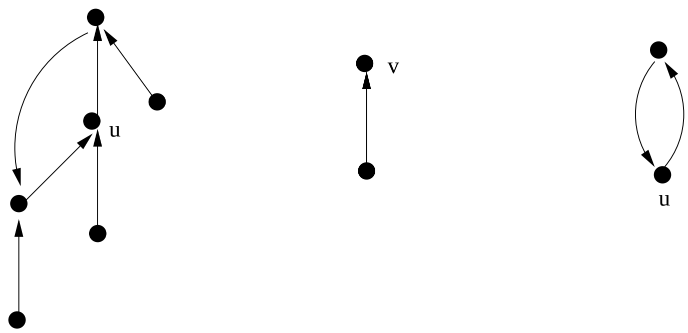

Figure 5: Examples of the graph F. One possible solution for the graph on the left (right) is to circle the vertex u and mark all other vertices cross. If vertex v in the middle graph is not in C then neither vertex is circled.

Theorem 6 *When the ring sizes n gets large, namely, $n \gg K$, the running time of our algorithm $\mathcal{A}_{MR}$ is $(1+o(1))(n/\sqrt{2})^K \cdot P \cdot Q$, where P is the time needed for topological sort of K sequences of length n, and Q is a polynomial in K.*

Proof: It is easy to see that the overall running time is the product of the running times of the three phases and P, the time for each topological sort. From Lemmas 2, 4 and 5, the overall running time is,

$$(1+o(1))\frac{n^{|L|+|N|}}{2^{(|L|+|N|)(c-2)}} \cdot (poly(K) + (2n)^{|D|}) \cdot (poly(K)(n/\sqrt{2})^{|C|}) \cdot P.$$

We have,

$$\frac{n^{|L|+|N|}}{2^{(|L|+|N|)(c-2)}} \cdot (2n)^{|D|} \cdot (n/\sqrt{2})^{|C|} \quad \leq \quad \frac{n^K}{2^{K/2}} \cdot \frac{1}{2^{(|L|+|N|)(c-2.5)}2^{-3|D|/2}}.$$

When the ring size n gets large, the value of δ is large and hence the size of the dominating set D approaches a small constant. When $c > 2.5$, the exponent of the second term in the above denominator is positive. Hence, the above expression approaches $(n/\sqrt{2})^K$. Therefore, the overall running time of algorithm $\mathcal{A}_{MR}$ is $(1+o(1))(n/\sqrt{2})^K \cdot P \cdot Q$. □

We note that the naive algorithm that enumerates all $2n$ possibilities for each ring takes $(2n)^K \cdot P$ time. Our algorithm essentially improves the term $(2n)^K$ to $(n/\sqrt{2})^K$.

Our algorithm achieves better running time for two important subclasses of inputs. Consider the subclass of *sparse* inputs: in the intersection graph of the rings, H, all the vertices are of low-degree. Thus, our algorithm terminates after Phase 1. The following comes directly from Lemma 2.

Corollary 7 *For any $c \geq 3$, if the maximal degree in H is smaller than $\log n/c$ then the running time of the algorithm is at most $(1+o(1))(\frac{n}{2^{c-2}})^K$. In particular, if the maximal degree in H is some constant $d \geq 1$ then the running time of $\mathcal{A}_{MR}$ is $(1+o(1))(n^{1-\frac{1}{d}})^K$.*

Consider now the subclass of *dense* inputs, where each node in the network appears in at least m rings, for some $m \geq 2$; then, as m grows larger, the running time of the algorithm approaches $(n/2)^K$. Formally,

Corollary 8 *If each node in the network appears in at least m rings, for some $m \geq 2$, then the running time of the algorithm is at most* $(1 + o(1))(\frac{n}{2^{(1-\frac{1}{m})}})^K$.

Proof: The running times of Phases 1 and 2 remain unchanged, therefore it is sufficient to show that the running time of Phase 3 reduces to $poly(K)(2 + n/2)^{|C|} \cdot 2^{|C|/m}$. The proof is similar to the proof of Lemma 5, except that the algorithm which implements Phase 3 is slightly changed. We describe the modified algorithm below.

Generally, when the input is *dense*, we get in Phase 3 that the $|C|$ remaining rings can be grouped to *clusters* of size at least m, such that in each cluster we need to try the two possible openings of a *single* ring. This enables one to reduce the number of clusters from $|C|/2$ (for arbitrary input, where $m = 2$), to $|C|/m$. In the following we explain how the clusters are constructed. Given a 'good' guess of the two last orderings for each ring $R_u \in C$ (out of which one is a correct opening), we construct the graph F, where each vertex corresponds to a vertex in C. Now, for any ring R_u, since each node in R_u appears in m rings in the network, we add a directed *solid* edge from R_u to some vertex v such that R_v contains the pivot of R_u, and at least $(m - 1)$ directed *dotted* edges from u to vertices representing other rings which contain the pivot of R_u.

Now, for any vertex $u \in F$ that is unmarked, we follow the solid edges on a directed path starting in u, until we have reached

(i) a marked vertex (either with a circle or with a cross), or

(ii) a vertex z from which there is a solid edge (z, u).

In both cases, we mark a cross on every unmarked vertex along the path (starting with u). In case (*ii*) we circle z. Also, in this case, we mark a cross on any vertex w such that there is a dotted edge (z, w). Finally, whenever we complete handling a path, we go over the directed graph F. For any vertex y, let (y, s) be the solid outgoing edge of y. If s is unmarked, and there exists a dotted edge (y, x), such that x is marked, we make the edge (y, x) solid, and (y, s) becomes dotted. (If there are a few dotted outgoing edges from y, that reach marked vertices, we choose one arbitrarily to become solid).

Now, we show that at most $|C|/m$ vertices are circled in F. Consider a dotted edge (z, w), and suppose that z is circled by the algorithm. We note that if z has less than m outgoing edges, then some ring that was opened in Phase 1 or 2 contains the pivot of z. This means that the opening of z is implied, and we can omit z from F. Thus, we may assume that z has at least m outgoing edges in F. Now, we argue that at the time that z is circled, we mark a cross on m unmarked vertices, i.e., all of the m outgoing edges of z reach unmarked vertices. Assume by contradiction that a vertex a is circled at some earlier time, and a vertex w having the dotted edges (a, w) and (z, w) is marked by a. Let w be the first of the m neighbors of z to be marked. By the algorithm, after we circle a, we go over the graph F and make dotted edges solid, if they lead to marked vertices. Therefore, the dotted edge (z, w) becomes solid. This implies that z cannot be marked by a circle, a contradiction.

By the above discussion, we get that the sets of m vertices marked by a cross by each circled vertex are *disjoint*.

Recall that we try the two openings only for circled vertices, and the openings for all other vertices in F are implied. Therefore, the running time of Phase 3 is reduced to $poly(K)(2 + \frac{n}{2})^{|C|} \cdot 2^{\frac{|C|}{m}}$. This completes the proof. □

Ring Clearance. In the ring clearance problem, we need to "clear" R_1 and reroute all the traffic through the other rings. In order for such transition to occur, it is assumed that R_1 intersects with each of the other rings. In other words, in the intersection graph H every vertex is a neighbor of the vertex w corresponding to R_1. Hence, $\{w\}$ is a dominating set for all vertices in H. We only need to apply Phase 3 of our algorithm. Using the simple analysis in Lemma 3, it is easy to see that any opening of R_1 limits the number of openings of any other ring to at most n. If we follow the more sophisticated pairing argument in Lemma 3, we only need to try a total of $(n/\sqrt{2})^K$ possibilities.

Corollary 9 *The algorithm solves the ring clearance problem in at most $(n/\sqrt{2})^K \cdot P \cdot Q$ steps, for rings of any length $n \geq 2$.*

Rings of distinct lengths. The analysis for the case where each ring R_u has a distinct size, n_u, is similar. We remove the singleton nodes and create the intersection graph H as before. For Phase 1, we say that a vertex u has low degree if it has fewer than $\delta_u = \log n_u / c$ neighbors. The running time of Phase 1 is at most $(1 + o(1)) \prod_{u \in N}(2n_u) \prod_{u \in L} \delta_u$. Similar to Lemma 2 we deduce,

$$\prod_{u \in N}(2n_u) \prod_{u \in L} \delta_u \leq \frac{\prod_{u \in L \cup N} n_u}{2^{(|L|+|N|)(c-2)}}.$$

In Phase 2, we find again a dominating set D, and we can bound $|D|$ by $|V - L - N| \cdot \frac{1+\ln(1+\delta)}{1+\delta}$, where $\delta = \min_u \delta_u$. When all ring sizes n_u get large, the size of D approaches a small constant. The running time for Phase 2 is $poly(K) + \prod_{u \in D}(2n_u)$. Finally, for Phase 3 we use the pairing algorithm as described in Lemma 5 for the vertices in C, and the running time is $poly(K) \prod_{u \in C}(n_u/\sqrt{2})$.

Theorem 10 *For rings with distinct sizes, when the ring sizes get large, the running time of algorithm $\mathcal{A}_{MR}$ is $(1 + o(1)) \prod_u (n_u/\sqrt{2}) \cdot P \cdot Q$, where the $o(1)$ term is a function of the ring sizes and c, P is the time needed for topological sort of K sequences, and Q is a polynomial in K.*

4 Relation to Other Problems

We now discuss how MRP relates to the shortest common supersequence (SCS) and feedback arc set (FAS) problems.

4.1 Shortest Common Supersequence

In SCS we are given K strings, $S = \{S_1, \ldots, S_K\}$, of lengths $n_1, \ldots, n_K$, over an alphabet Σ, where $|\Sigma| = \mathcal{N}$. W.l.o.g., we assume that every letter in Σ appears in the input. We seek a supersequence T for S of minimum length. MRP defines the following natural variant of SCS. A *two-way cyclic* permutation of a string allows cyclic shifts of the string in the forward and reverse directions. For example, the string *abcd* has 4 forward shifts, *abcd*, *bcda*, *cdab* and *dabc*, and 4 reverse shifts, *dcba*, *cbad*, *badc* and *adcb*. In the *two-way cyclic SCS (2Cyclic-SCS)* problem, we seek a string T of minimum length, such that there exists a two-way cyclic permutation of each string $S_1, \ldots, S_K$ in S that is a subsequence of T. We say that T is a *2cyclic supersequence* for S. A supersequence T of length $\mathcal{N}$ corresponds to a master ring for the set of rings defined by $S_1, \ldots, S_K$. We show in Section 5 that 2Cyclic-SCS is NP-hard.

The SCS problem is known to be hard to approximate. In particular, Jiang and Li [10] showed that there exists a constant $\varepsilon > 0$ such that if SCS has a polynomial time approximation algorithm with ratio $\log^{\varepsilon} K$, then NP is contained in DTIME$(2^{polylog(K)})$. The best known approximation ratio is $\frac{K+3}{4}$, due to Fraser and Irving [7]. Middendorf considered in [11] a number of variants of SCS. This includes the Cyclic-SCS problem, in which the strings in S can be cyclically permuted in the *same* direction. The paper shows that this problem is NP-hard. (Cyclic-SCS solves MRP in the case where each ring has a *fixed* orientation.) On the other hand, Permutation-SCS, where each string S_k can be permuted to any one of the $n_k!$ possibilities, is shown in [11] to be polynomially solvable for strings of any length. This implies that MRP can be solved in polynomial time for inputs where $n_k \leq 3$, for $1 \leq k \leq K$.

The 2Cyclic-SCS problem can be optimally solved using the following natural dynamic programming algorithm. Construct the supersequence T using a configuration vector $((s_1, \ell_1, o_1), \ldots, (s_K, \ell_K, o_K), last)$. The triple (s_k, ℓ_k, o_k) indicates that the subsequence of T corresponding to S_k starts with the letter s_k, ℓ_k is the number of letters in S_k covered by T (so far), and o_k is the 'orientation' of S_k in T (i.e., forward or reverse shift); *last* is the last letter added to T.

We can find a supersequence T of optimal length using the following recursion.

$$\ell((s_1, \ell_1, o_1), \ldots, (s_K, \ell_K, o_K), v) = 1 + \min_{u \in \Sigma} \ell((s_1, \ell'_1, o_1), \ldots, (s_K, \ell'_K, o_K), u),$$

where $\ell'_k = \ell_k - 1$ if v is the ℓ_k-th letter in S_k (when we take a cyclic shift of S_k in direction o_k, and the letter starting the string is s_k); otherwise $\ell'_k = \ell_k$. We solve the 2Cyclic-SCS problem by finding

$$\min_{\substack{s_1, \ldots, s_K \\ o_k \in \{left, right\} \\ last \in \Sigma}} \ell((s_1, n_1, o_1), \ldots, (s_K, n_K, o_K), last).$$

Thus, the natural dynamic programming algorithm yields an optimal solution for 2Cyclic-SCS in $O(\mathcal{N} 2^K \prod_{k=1}^{K} n_k^2)$ steps. Note that the decision version of 2Cyclic-SCS is NP-hard, already in the case where we seek a supersequence of length $\mathcal{N}$, the total number of distinct letters in the input. This follows directly from the hardness of MRP. We can solve this problem by applying algorithm $\mathcal{A}_{MR}$.

4.2 Feedback Arc Set

MRP relates also to the *feedback arc set (FAS)* problem in directed graphs, which is known to be NP-hard [8]. Consider the special case of MRP in which the orientation for each of the rings is given. We denote this *oriented* version $MRP_{\mathcal{O}}$. We can view $MRP_{\mathcal{O}}$ as the following variant of FAS, that we call *exact subset FAS*. We have a directed graph $G = (V, E)$, and a set of K (directed) cycles in G, $R = \{R_1, \ldots, R_K\}$. Let $G' = (V', E')$ be the subgraph induced by the vertices and edges in R. We seek a subset of K edges in E' whose deletion leaves G' acyclic, such that in each of the cycles $R_1, \ldots, R_K$ we omit *exactly* one edge. Such a subset of vertices exists iff we have a solution for the corresponding $MRP_{\mathcal{O}}$ instance. Since we are given the orientation for each of the rings, we can apply only Phase 3 of algorithm $\mathcal{A}_{MR}$. By finding a master ring, we solve the *exact subset FAS* problem. Hence, we have

Corollary 11 *For any $K \geq 1$ and $n_k \geq 2$, for all $1 \leq k \leq K$, exact subset FAS on the subgraph G' induced by K cycles of the lengths $n_1, \ldots, n_K$ can be solved in $P \cdot (\prod_{k=1}^{K} n_k / (\sqrt{2})^K)$ steps, where P is a polynomial of K.*

5 Hardness Results

5.1 Hardness of MRP

Theorem 12 *The master ring problem is NP-complete.*

Proof: The proof is by reduction from *One-in-Three 3SAT*.
One-in-Three 3SAT (e.g., see [8])
Input: A $p \times q$ matrix A such that (i) each entry is 0 or 1, and (ii) every row contains exactly three 1s.
Question: Is there a vector $z \in \{0,1\}^q$ satisfying $Az = 1_p$? (The vector 1_p denotes the p-dimensional all one vector.)

A key gadget in our hardness proof is a set of three rings $\{gde,\ gef,\ gfd\}$, that have exactly three master rings: $dgef$, $degf$ and $defg$.

We associate these three master rings with three 0-1 valued solutions of an equality $z_i + z_j + z_k = 1$ in a given One-in-Three 3SAT instance.

In the following, we assume that $q \geq 3$, and we denote the index sets of rows and columns of A by $\mathrm{row}(A) = \{1, 2, \ldots, p\}$ and $\mathrm{col}(A) = \{1, 2, \ldots, q\}$, respectively. Given an instance, a $p \times q$ matrix A, of One-in-Three 3SAT, we construct an instance of the master ring problem with $(2 + 3p)$ rings and $(4p + 3q)$ nodes as follows. First, we introduce a set of $(4p + 3q)$ nodes defined by;

$$\{a_i \mid i \in \mathrm{col}(A)\} \cup \{b_i \mid i \in \mathrm{col}(A)\} \cup \{c_i \mid i \in \mathrm{col}(A)\} \\ \cup \{d_h \mid h \in \mathrm{row}(A)\} \cup \{e_h \mid h \in \mathrm{row}(A)\} \cup \{f_h \mid h \in \mathrm{row}(A)\} \cup \{g_h \mid h \in \mathrm{row}(A)\}.$$

Let R_a and R_b be a pair of rings corresponding to $c_1 a_1 c_2 a_2 \cdots c_q a_q$ and $c_1 b_1 c_2 b_2 \cdots c_q b_q$, respectively. It is easy to verify that any master ring R of the two rings $\{R_a, R_b\}$ satisfies (i) $c_1 c_2 \cdots c_q$ is a subring of R and (ii) for any column-index $i \in \mathrm{col}(A)$, R contains exactly one of the following two subrings; $c_i a_i b_i c_{i+1}$ and $c_i b_i a_i c_{i+1}$ where we identify c_{q+1} with c_1. For each row-index $h \in \mathrm{row}(A)$, we introduce three rings S_h, T_h, U_h, defined as follows. Since each row of A has three 1s, there exist three column-indices $i, j, k \in \mathrm{col}(A)$ satisfying $i < j < k$ and $a_{hi} = a_{hj} = a_{hk} = 1$. We introduce three rings S_h, T_h, U_h defined by the following corresponding sequences;

$$S_h : g_h d_h c_i a_i b_i e_h; \quad T_h : g_h e_h c_j a_j b_j f_h; \quad \text{and} \quad U_h : g_h f_h c_k a_k b_k d_h.$$

Here, we define an instance of the master ring problem with the set of $(2+3p)$ rings $\{R_a, R_b\} \cup \{S_h \mid h \in \mathrm{row}(A)\} \cup \{T_h \mid h \in \mathrm{row}(A)\} \cup \{U_h \mid h \in \mathrm{row}(A)\}$ described above.

Next, we show that if there exists a master ring R of the set of $(2+3p)$ rings described above, then the equality system $Az = 1_p$ has a 0-1 valued solution. We define a vector $z' \in \{0,1\}^q$ by setting $z'_i = 0$ if $c_i a_i b_i c_{i+1}$ is a subsequence of R and $z'_i = 1$ if $c_i b_i a_i c_{i+1}$ is a subsequence of R. Let $h \in \mathrm{row}(A)$ be any row-index of the matrix A, and $i, j, k \in \mathrm{col}(A)$ be column-indices satisfying $i < j < k$ and $a_{hi} = a_{hj} = a_{hk} = 1$. Then, we only need to show that $z'_i + z'_j + z'_k = 1$. The existence of 4 nodes $\{d_h, e_h, f_h, g_h\}$ implies that R has exactly one of the three rings R_1, R_2, R_3 as a subring where R_1, R_2, R_3 are defined by the following corresponding sequences; $R_1 : d_h g_h e_h f_h$; $R_2 : d_h e_h g_h f_h$; $R_3 : d_h e_h f_h g_h$.

Case (1) When R_1 is a subring of R. The set of three rings $\{R_1, T_h, U_h\}$ has a unique (node-wise minimal) master ring R' corresponding to $d_h g_h e_h c_j a_j b_j f_h c_k a_k b_k$. The uniqueness implies that R' is a subring of R. Since R is a master ring of $\{R', S_h, R_a, R_b\}$, it is easy to see that R has a subring R'' corresponding to $d_h g_h e_h b_i a_i c_j a_j b_j f_h c_k a_k b_k c_i$ and thus R contains three sequences $c_i b_i a_i c_{i+1}$, $c_j a_j b_j c_{j+1}$ and $c_k a_k b_k c_{k+1}$. From the above, we obtain $(z'_i, z'_j, z'_k) = (1,0,0)$ and $z'_i + z'_j + z'_k = 1$.

Case (2) When R_2 is a subring of R, we can show that $(z'_i, z'_j, z'_k) = (0,1,0)$, similar to Case (1).

Case (3) When R_3 is a subring of R, we can show that $(z'_i, z'_j, z'_k) = (0,0,1)$, similar to Case (1).

Lastly, we show the inverse implication that if the equality system $Az = 1_p$ has a 0-1 valued solution $z' \in \{0,1\}^q$, then a master ring R exists. Here we introduce $(1+q)$ rings. First, we define a ring $R(z')$ corresponding to $c_1 \alpha_1 \beta_1 c_2 \alpha_2 \beta_2 \cdots c_q \alpha_q \beta_q$, where for each column-index $i \in \mathrm{col}(A)$, $z'_i = 0$ implies $(\alpha_i, \beta_i) = (a_i, b_i)$, and $z'_i = 1$ implies $(\alpha_i, \beta_i) = (b_i, a_i)$. It is clear that both R_a and R_b are subrings of $R(z')$. For each row-index $h \in \mathrm{row}(A)$, we introduce a ring $R_h(z')$ as follows. Let $i, j, k \in \mathrm{col}(A)$ be column-indices satisfying $i < j < k$ and $a_{hi} = a_{hj} = a_{hk} = 1$. We set

$$R_h(z') \text{ to a ring corresponding to } \begin{cases} d_h g_h e_h b_i a_i c_j a_j b_j f_h c_k a_k b_k c_i, & \text{if } (z'_i, z'_j, z'_k) = (1,0,0), \\ e_h g_h f_h b_j a_j c_k a_k b_k d_h c_i a_i b_i c_j, & \text{if } (z'_i, z'_j, z'_k) = (0,1,0), \\ f_h g_h d_h b_k a_k c_i a_i b_i e_h c_j a_j b_j c_k, & \text{if } (z'_i, z'_j, z'_k) = (0,0,1). \end{cases}$$

In any case, S_h, T_h and U_h are subrings of $R_h(z')$. Then, it is sufficient to show that there exists a master ring R of the set of $(1+p)$ rings $\{R(z')\} \cup \{R_h(z') \mid h \in \mathrm{row}(A)\}$ defined above. Each node in $N' = \{d_h \mid h \in \mathrm{row}(A)\} \cup \{e_h \mid h \in \mathrm{row}(A)\} \cup \{f_h \mid h \in \mathrm{row}(A)\} \cup \{g_h \mid h \in \mathrm{row}(A)\}$ appears in exactly one ring in $\{R(z')\} \cup \{R_h(z') \mid h \in \mathrm{row}(A)\}$, and thus we do not need to consider the nodes N'. It is obvious that for each row-index $h \in \mathrm{row}(A)$, the ring obtained from $R_h(z')$ by deleting nodes in N' is a subring of $R(z')$. From the above, the existence of a master ring R of $\{R(z')\} \cup \{R_h(z') \mid h \in \text{row}(A)\}$ is now trivial. □

Corollary 13 *The master ring problem remains NP-complete even if a given set of rings satisfies the condition that every pair of intersecting rings have at least two nodes in common.*

The claim follows from the fact that the instance of the master ring problem defined in the above proof satisfies the condition.

5.2 Hardness of 2Cyclic-SCS

Formally, the 2Cyclic-SCS problem is defined as follows.

Input: A set of K strings, $S = \{S_1, \ldots, S_K\}$ over an alphabet Σ, where $|\Sigma| = \mathcal{N}$, and an integer $m \geq 1$.

Question: Is there a 2cyclic supersequence T for S of length at most m?

In the following we give a proof of hardness for the problem.

Theorem 14 *Given m, it is NP-hard to determine if 2Cyclic-SCS has a solution of length at most m.*

Proof: The proof is by reduction from *Vertex Cover (VC)*.

Input: An undirected graph $G = (V, E)$ and a positive integer $\ell \leq |V|$.

Question: Is there a subset of vertices $V' \subseteq V$ of size at most ℓ, such that for each edge $(v_i, v_j) \in E$ at least one of v_i and v_j is in V'?

Given an instance I of Vertex Cover with $|V| = n$, we construct the following instance I' of 2Cyclic-SCS. For each edge $(v_i, v_j) \in E$, we define the set of strings $\{x_p v_i v_j y_p,\ x_p v_j v_i y_p \mid 1 \leq p \leq n+1\}$; that is, $K = 2(n+1)|E|$. Now, we show that there is a VC of size at most ℓ for I iff there is a 2cyclic supersequence T for I' of length $L(T) \leq n + 2(n+1) + \ell$.

(i) Suppose that there is a VC of size at most ℓ, $v_{i1}, \ldots, v_{i\ell}$; denote by $v_{j1}, \ldots, v_{j(n-\ell)}$ the set of remaining vertices in V. Then the following is a 2cyclic supersequence for I' of length $3n + 2 + \ell$.

$$x_1 \cdots x_{n+1} v_{i1} \cdots v_{i\ell} v_{j1} \cdots v_{j(n-\ell)} v_{i1} \cdots v_{i\ell} y_1 \cdots y_{n+1}.$$

(ii) Now, suppose that we have a supersequence T with $L(T) \leq 3n + 2 + \ell$, then we show that there is a VC of size at most ℓ. Assume towards contradiction that there exists an edge $(v_i, v_j) \in E$, such that both v_i and v_j appear in T exactly once. We need to handle two cases.

  (a) The vertex v_i appears in T before v_j. We further distinguish between three cases. If x_p appears before v_i then, by considering all possible two-way cyclic permutations of $x_p v_i v_j y_p$ and $x_p v_j v_i y_p$, we get that either x_p or y_p has to appear twice in T. Similarly, in the case where x_p appears is T between v_i and v_j, or if x_p appears after v_j. In all these cases, we get that x_p or y_p appear twice in T. This implies that $L(T) \geq n + 3(n+1)$. A contradiction.

From the above discussion, by taking the set of 'repeated' vertices in T, each edge $(v_i, v_j) \in E$ has at least one end in the set, thus we get a vertex cover for G. Since the size of the alphabet for I' is $(n + 2(n+1))$, we get that the size of the VC is at most $3n + 2 + \ell - (n + 2(n+1)) = \ell$.

(b) The vertex v_j appears in T before v_i. The proof is similar; we omit the details.

□

6 Open Problems

Consider the following parameterized version of the Permutation-SCS. Each string S_k, $1 \leq k \leq K$, is associated with a subset of permutations, Π_k, and we seek a supersequence T of minimum length, such that there exists a permutation of S_k in Π_k that is a subsequence of T. We call this problem *Perm-SCS*(Π_k). Indeed, the Cyclic-SCS problem is a special case of this problem, in which Π_k is the set of n_k cyclic shifts of S_k, in a single direction. As shown in [11], this special case of the problem is NP-hard. We have shown (in Theorem 14) that if we extend the permutation sets in the Cyclic-SCS, so that Π_k is the set of cyclic shifts in *two* directions (2Cyclic-SCS), the problem remains NP-hard. On the other hand, when Π_k is the set of *all* possible permutations of S_k (Permutation-SCS), the problem is solvable in polynomial time [11]. Determining whether Perm-SCS(Π_k) is polynomially solvable on other classes of inputs remains an open problem.

Finally, a natural variant of MRP which is of practical interest, is to identify a *maximum* subset of rings for which we can find a master ring, in any given network.

Acknowledgments

We thank four anonymous referees for many insightful comments on the paper.

References

[1] S. Acharya, B. Gupta, P. Risbood, A. Srivastava. *Hitless Network Engineering of SONET Rings*, Globecom 2003.

[2] S. Acharya, B. Gupta, P. Risbood, A. Srivastava. *In-service Optimization of stacked SONET Rings.* Manuscript, 2004.

[3] N. Alon and J. H. Spencer. The Probabilistic Method, Second Edition. Wiley-Interscience, 2000.

[4] T. H. Cormen, C. E. Leiserson, R. L. Rivest and C. Stein, *Introduction to Algorithms*, 2nd Edition, MIT Press and McGraw-Hill, 2002.

[5] D. E. Foulser, M. Li and Q. Yang, *A Theory of Plan Merging*, Artificial Intelligence, 57, 1992, pp. 143–181.

[6] C. B. Fraser, *subsequences and Supersequences of Strings.* Ph.D. Thesis, Dept. of Computer Science, University of Glasgow, 1995.

[7] C. B. Fraser and R. W. Irving , *Approximation algorithms for the shortest common supersequence*, Nordic J. Comp. 2, 1995, pp. 303–325.

[8] M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W.H. Freeman, 1979.

[9] S.Y. Itoga, *The String Merging Problem*, BIT, 21, 1981, pp.20–30.

[10] T. Jiang and M. Li, *On the Approximation of Shortest Common Supersequences and Longest Common Subsequences*, SIAM Journal on Computing, 24(5), October 1995, pp. 1122–1139.

[11] M. Middendorf, More on the complexity of common superstring and supersequence problems, *Theoretical Computer Science* 125 (1994), 205-228.

[12] Mobius network management and optimization systems. Lucent Technologies Proprietary. Internal website: http://www-zoo.research.bell-labs.com/∼mobius/.

[13] R. Ramaswami and K. Sivarajan. *Optical networks: a practical perspective.* (Morgan Kaufmann Publishers Inc., San Francisco, 1998).

[14] H. Shachnai and L. Zhang The Master Ring Problem In *Proc. of 2005 Int. Conference on Analysis of Algorithms*, pp. 287–296.