

APPROXIMATING THE MINIMUM EQUIVALENT DIGRAPH*

SAMIR KHULLER[†], BALAJI RAGHAVACHARI[‡], AND NEAL YOUNG[§]

Abstract. The minimum equivalent graph (MEG) problem is as follows: given a directed graph, find a smallest subset of the edges that maintains all reachability relations between nodes. This problem is NP-hard; this paper gives an approximation algorithm achieving a performance guarantee of about 1.64 in polynomial time. The algorithm achieves a performance guarantee of 1.75 in the time required for transitive closure.

The heart of the MEG problem is the minimum strongly connected spanning subgraph (SCSS) problem—the MEG problem restricted to strongly connected digraphs. For the minimum SCSS problem, the paper gives a practical, nearly linear-time implementation achieving a performance guarantee of 1.75.

The algorithm and its analysis are based on the simple idea of contracting long cycles. The analysis applies directly to 2-EXCHANGE, a general “local improvement” algorithm, showing that its performance guarantee is 1.75.

Key words. directed graph, approximation algorithm, strong connectivity, local improvement

AMS subject classifications. 68R10, 90C27, 90C35, 05C85, 68Q20

1. Introduction. Connectivity is fundamental to the study of graphs and graph algorithms. Recently, many approximation algorithms for finding minimum subgraphs that meet given connectivity requirements have been developed [1], [9], [11], [15], [16], [24]. These results provide practical approximation algorithms for NP-hard network-design problems via an increased understanding of connectivity properties.

Until now, the techniques developed have been applicable only to *undirected* graphs. We consider a basic network-design problem in *directed* graphs [2], [12], [13], [18], which is as follows: given a digraph, find a smallest subset of the edges (forming a *minimum equivalent graph (MEG)*) that maintains all reachability relations of the original graph.

When the MEG problem is restricted to strongly connected graphs we call it *the minimum strongly connected spanning subgraph (SCSS) problem*. When the MEG problem is restricted to acyclic graphs we call it *the acyclic MEG problem*. The MEG problem reduces in linear time [5] to a single acyclic problem given by the so-called strong component graph, together with one minimum SCSS problem for each strong component (given by the subgraph induced by that component). Furthermore, the reduction preserves approximation in the sense that c -approximate solutions to the subproblems yield a c -approximate solution to the original problem.

Moyses and Thompson [18] observe this decomposition and give exponential-time algorithms for the restricted problems. Hsu [13] gives a polynomial-time algorithm for the acyclic MEG problem.

The related problem of finding a *transitive reduction* of a digraph—a smallest set of edges yielding the same reachability relations—is studied by Aho, Garey, and Ullman [2]. Transitive reduction differs from the MEG problem in that the edges in the transitive reduction are not required to be in the original graph. However, the transitive reduction problem decomposes just like the MEG problem into acyclic and strongly connected instances. For any strongly connected instance, a transitive reduction is given by any Hamilton cycle through the vertices.

*Received by the editors September 7, 1993; accepted for publication (in revised form) April 12, 1994.

[†]Computer Science Department and Institute for Advanced Computer Studies, University of Maryland, College Park, Maryland 20742 (samir@cs.umd.edu). This research was supported by National Science Foundation Research Initiation Award CCR-9307462.

[‡]Computer Science Department, University of Texas at Dallas, Richardson, Texas 75083-0688 (rbk@utdallas.edu).

[§]AT&T Bell Labs, Room 2D-145, 600 Mountain Avenue, Murray Hill, New Jersey 07974. Part of this work was done while at the University of Maryland Institute for Advanced Computer Studies and was supported in part by National Science Foundation grants CCR-8906949 and CCR-9111348 (ney@research.att.com).

For an acyclic instance, the transitive reduction is unique and, as Aho, Garey, and Ullman observe, equivalent to an MEG: it consists of those edges (u, v) for which there is no alternate path from u to v . In fact, Aho, Garey, and Ullman show that the transitive reduction problem is *equivalent* to the transitive closure problem. Thus, the acyclic MEG problem reduces to transitive closure.

The acyclic MEG problem can be solved in polynomial time, whereas the minimum SCSS problem is NP-hard [8], [20]. Consequently, this paper focuses on approximation algorithms for the minimum SCSS problem. By the observations of the preceding paragraphs, the performance guarantees obtained for the minimum SCSS problem carry over to the general MEG problem with the overhead of solving a single instance of transitive closure.

1.1. Our results. Given a strongly connected graph, our basic algorithm finds as long a cycle as it can, contracts the cycle, and recurses. The contracted graph remains strongly connected. When the graph finally collapses into a single vertex, the algorithm returns the set of edges contracted during the course of the algorithm as the desired SCSS.

The algorithm achieves a performance guarantee of any constant *greater* than $\pi^2/6 \approx 1.645$ in polynomial time. We give a nearly linear-time version that achieves a performance guarantee of 1.75. We give examples showing lower bounds on the performance guarantees of the algorithm. For the general algorithm, the lower bounds are slightly above 1.5. For the nearly linear-time version, the lower bound is 1.75, matching the upper bound.

The performance guarantee analysis extends directly to a simple “local improvement” algorithm called 2-EXCHANGE. 2-EXCHANGE starts with the given digraph and performs the following local improvement step as long as it is applicable: find two edges in the current graph that can be replaced by one edge from the original graph, maintaining strong connectivity. Similar local improvement algorithms are natural candidates for many optimization problems but often elude analysis. We prove that the performance guarantee of 2-EXCHANGE is 1.75.

A natural improvement to the cycle-contraction algorithm is to modify the algorithm to solve the problem optimally once the contracted graph has no cycles longer than a given length c . For instance, for $c = 3$ this modification improves the performance guarantee to $\pi^2/6 - 1/36 \approx 1.617$. We use SCSS_c to denote the minimum SCSS problem restricted to digraphs with no cycle longer than c . The minimum SCSS_2 problem is trivial. The minimum SCSS_3 problem can be solved in polynomial time as shown by Khuller, Raghavachari, and Young [14]. However, further improvement in this direction is limited: we show that the minimum SCSS_5 problem is NP-hard. In fact, we show that the minimum SCSS_{17} problem is MAX SNP-hard. This precludes the possibility of a polynomial-time approximation scheme, assuming $\text{P} \neq \text{NP}$ [4].

1.2. Other related work. The union of any incoming branching and any outgoing branching from the same root yields an SCSS with at most $2n - 2$ edges (where n is the number of vertices in the graph). This is a special case of the algorithm given by Frederickson and Jala [6] that uses minimum weight branchings [7] to achieve a performance guarantee of 2 for weighted graphs. Since any SCSS has at least n edges, this yields a performance guarantee of 2 for the SCSS problem.

Any *minimal* SCSS (one from which no edge can be deleted) has at most $2n - 2$ edges and yields a performance guarantee of 2. The problem of efficiently finding a minimal SCSS is studied by Simon [21]. Gibbons et al. [10] give a parallel algorithm.

A related problem in undirected graphs is to find a smallest subset of the edges forming a biconnected (respectively, bridge-connected (i.e., 2-edge-connected)) spanning subgraph of a given graph. These problems are NP-hard. Khuller and Vishkin [15] give a depth-first-search- (DFS-) based algorithm that achieves a factor of $\frac{5}{3}$ for biconnectivity and $\frac{3}{2}$ for bridge-connectivity. Garg, Santosh, and Singla [9] subsequently improve the approximation

factor for biconnectivity, using a similar approach, to $\frac{3}{2}$. None of these methods appear to extend directly to the minimum SCSS problem.

Undirected graphs having bounded cycle length have bounded tree width. Arnborg, Lagergren, and Seese [3] have shown that many NP-hard problems, including the minimum biconnected-spanning-subgraph problem, have polynomial-time algorithms when restricted to such graphs.

2. Preliminaries. To *contract* a pair of vertices u, v of a digraph is to replace u and v (and each occurrence of u or v in any edge) by a single new vertex and delete any subsequent self-loops and multiedges. Each edge in the resulting graph is identified with the corresponding edge in the original graph or, in the case of multiedges, the single remaining edge is identified with any one of the corresponding edges in the original graph. To contract an edge (u, v) is to contract the pair of vertices u and v . To contract a set S of pairs of vertices in a graph G is to contract the pairs in S in arbitrary order. The contracted graph is denoted by G/S . Contracting an edge is also analogously extended to contracting a set of edges.

Let $OPT(G)$ be the minimum size of any subset of the edges that strongly connects G . In general, the term cycle refers only to simple cycles.

3. Lower bounds on $OPT(G)$. We begin by showing that if a graph has no long cycles, then the size of any SCSS is large.

LEMMA 3.1 (Cycle lemma). *For any strongly connected directed graph G with n vertices, if a longest cycle of G has length C , then*

$$OPT(G) \geq \frac{C}{C-1}(n-1).$$

Proof. Starting with a minimum-size subset that strongly connects the graph, repeatedly contract cycles in the subset until no cycles are left. Observe that the maximum cycle length does not increase under contractions. Consequently, for each cycle contracted, the ratio of the number of edges contracted to the decrease in the number of vertices is at least $\frac{C}{C-1}$. Since the total decrease in the number of vertices is $n-1$, at least $\frac{C}{C-1}(n-1)$ edges are contracted. \square

Note that the above lemma gives a lower bound which is existentially tight. For all values of C , there exist graphs for which the bound given by the lemma is equal to $OPT(G)$. Also note that C has a trivial upper bound of n and, using this, we get a lower bound of n for $OPT(G)$, which is the known trivial lower bound.

LEMMA 3.2 (Contraction lemma). *For any directed graph G and set of edges S ,*

$$OPT(G) \geq OPT(G/S).$$

Proof. Any SCSS of G contracted around S (treating the edges of S as pairs) is an SCSS of G/S . \square

4. Cycle-contraction algorithm. The algorithm follows. Fix k as any positive integer.

CONTRACT-CYCLES $_k(G)$

- 1 **for** $i = k, k-1, k-2, \dots, 2$
- 2 **while** the graph contains a cycle with at least i edges
- 3 Contract the edges on such a cycle.
- 4 **return** the contracted edges

In §6, we will show that the algorithm can be implemented to run in $O(m\alpha(m, n))$ time for the case $k = 3$ and in polynomial time for any fixed value of k , where m is the number

of edges. It is clear that the edge set returned by the algorithm strongly connects the graph. The following theorem establishes an upper bound on the number of edges returned by the algorithm.

THEOREM 4.1. *CONTRACT-CYCLES_k(G) returns at most $c_k \cdot OPT(G)$ edges, where*

$$\frac{\pi^2}{6} \leq c_k \leq \frac{\pi^2}{6} + \frac{1}{(k-1)k}.$$

Proof. Initially, let the graph have n vertices. Let n_i vertices remain in the contracted graph after contracting cycles with i or more edges ($i = k, k-1, \dots, 2$).

How many edges are returned? In contracting cycles with at least k edges, at most $\frac{k}{k-1}(n - n_k)$ edges are contributed to the solution. For $i < k$, in contracting cycles with i edges, $\frac{i}{i-1}(n_{i+1} - n_i)$ edges are contributed. Thus the number of edges returned is at most

$$\frac{k}{k-1}(n - n_k) + \sum_{i=2}^{k-1} \frac{i}{i-1}(n_{i+1} - n_i) \leq \left(1 + \frac{1}{k-1}\right)n + \sum_{i=3}^k \frac{n_i - 1}{(i-1)(i-2)}.$$

Clearly $OPT(G) \geq n$. For $2 \leq i \leq k$, when n_i vertices remain, no cycle has more than $i - 1$ edges. By Lemmas 3.1 and 3.2, $OPT(G) \geq \frac{i-1}{i-2}(n_i - 1)$. Thus the number of edges returned, divided by $OPT(G)$, is at most

$$\frac{\left(1 + \frac{1}{k-1}\right)n}{OPT(G)} + \sum_{i=3}^k \frac{\frac{n_i-1}{(i-1)(i-2)}}{OPT(G)} \leq \frac{\left(1 + \frac{1}{k-1}\right)n}{n} + \sum_{i=3}^k \frac{\frac{n_i-1}{(i-1)(i-2)}}{\frac{i-1}{i-2}(n_i - 1)} = \frac{1}{k-1} + \sum_{i=1}^{k-1} \frac{1}{i^2} = c_k.$$

Using the identity (from [17, p. 75]) $\sum_{i=1}^{\infty} 1/i^2 = \pi^2/6$, we get

$$\begin{aligned} \frac{\pi^2}{6} \leq c_k &= \frac{\pi^2}{6} + \frac{1}{k-1} - \sum_{i=k}^{\infty} \frac{1}{i^2} \\ &\leq \frac{\pi^2}{6} + \frac{1}{k-1} - \sum_{i=k}^{\infty} \frac{1}{i(i+1)} \\ &= \frac{\pi^2}{6} + \frac{1}{k-1} - \frac{1}{k} \\ &= \frac{\pi^2}{6} + \frac{1}{(k-1)k}. \quad \square \end{aligned}$$

If desired, standard techniques can yield more accurate estimates of c_k , e.g., $c_k = \pi^2/6 + 1/2k^2 + O(1/k^3)$. If the graph initially has no cycle longer than ℓ ($\ell \geq k$), then the analysis can be generalized to show a performance guarantee of $(k^{-1} - \ell^{-1})/(1 - k^{-1}) + \sum_{i=1}^{k-1} 1/i^2$. For instance, in a graph with no cycle longer than 5, the analysis bounds the performance guarantee (when $k = 5$) by 1.424.

Table 1 gives lower and upper bounds on the performance guarantee of the algorithm for small values of k and in the limit as $k \rightarrow \infty$. The lower bounds are shown in the next section.

4.1. Lower bounds on the performance ratio. In this section, we present lower bounds on the performance ratio of $CONTRACT-CYCLES_k(G)$. The graph in Fig. 1 has $\frac{n}{2k-2}$ groups of vertices. Each group consists of a $(2k - 2)$ -cycle “threaded” with a k -cycle.

In the first iteration, $CONTRACT-CYCLES_k(G)$ can contract the k -cycle within each group, leaving the graph with only 2-cycles. The algorithm subsequently must contract all the remaining edges. Thus, all the $(3k - 2)\frac{n}{2k-2} - 2$ edges are in the returned SCSS. The graph contains a Hamilton cycle and the optimal solution is thus n . Hence, for arbitrarily large n , $1 + \frac{k}{2k-2} - 2/n$ is a lower bound on the performance guarantee of $CONTRACT-CYCLES_k(G)$. As k approaches ∞ , the lower bound tends to 1.5.

TABLE I
Bounds on the performance guarantee.

k	Upper bound	Lower bound
3	1.750	1.750
4	1.694	1.666
5	1.674	1.625
∞	1.645	1.500




FIG. 1. *Bad example for CONTRACT-CYCLES_k(G).*

5. 2-EXCHANGE algorithm. In this section, we use the cycle-contraction analysis to show that 2-EXCHANGE has a performance guarantee of 1.75. 2-EXCHANGE is a special case of k -EXCHANGE, which is defined as follows.

k -EXCHANGE($G = (V, E)$) — *Local improvement algorithm.*

- 1 $E' \leftarrow E$
- 2 **while** the following improvement step is possible
- 3 Pick a set E_k of k edges in E' and a set E_{k-1} of up to $k - 1$ edges in E such that the set of edges $E'' = (E' - E_k) \cup E_{k-1}$ forms an SCSS.
- 4 $E' \leftarrow E''$.
- 5 **return** E'

Note that for fixed k , each step can be performed in polynomial time and reduces the size of E' , so k -EXCHANGE runs in polynomial time. The following theorem shows that the approximation factor achieved by 2-EXCHANGE is 1.75.

THEOREM 5.1. *The performance guarantee of 2-EXCHANGE is 1.75.*

Proof. We will show that the edges output by 2-EXCHANGE(G) could be output by CONTRACT-CYCLES₃(G). Thus, the performance guarantee of 1.75 for CONTRACT-CYCLES₃ carries over to 2-EXCHANGE.

First we show that the performance guarantee is at most 1.75. Let E' be the set of edges returned by 2-EXCHANGE($G = (V, E)$). Run CONTRACT-CYCLES₃ on the graph $G' = (V, E')$. Let H be the set of edges contracted during the first iteration when cycles of at least three edges are contracted. The resulting graph G'/H is strongly connected and has only 2-cycles. Such a graph has a tree-like structure. In particular, an edge (u, v) is present if and only if the reverse edge (v, u) is present.




FIG. 2. Worst-case example for 2-EXCHANGE.

The important observation is that G/H is equivalent to G'/H . Clearly G'/H is a subgraph of G/H ; to prove the converse, suppose that some edge (u, v) of G/H was not in G'/H . Consider adding edge (u, v) to G'/H . By the structure of G'/H , u and v are not adjacent in G'/H , and for each edge on the path from v to u , the reverse edge is also in G'/H . If (u, v) is added to G'/H , these (at least two) reverse edges can be deleted from G'/H without destroying the strong connectivity of G'/H . Consequently, the original edge in G corresponding to (u, v) can be added to G' , and the original edges in G' corresponding to the reverse edges can be deleted from G' , without destroying the strong connectivity of G' . This contradicts the fact that E' was output by 2-EXCHANGE(G), since E' is eligible for an improvement step.

Next consider executing CONTRACT-CYCLES₃(G). Since G/H is equivalent to G'/H , the sequence of cycles chosen in the first iteration of CONTRACT-CYCLES₃(G') could also be chosen by the first iteration of CONTRACT-CYCLES₃(G). Similarly, the second iteration in CONTRACT-CYCLES₃(G') could be mimicked by CONTRACT-CYCLES₃(G), in which case CONTRACT-CYCLES₃(G) would return the same edge set as CONTRACT-CYCLES₃(G'). Since E' is minimal (otherwise an improvement step applies), the edge set returned is exactly E' . Thus, the upper bound on the performance guarantee of CONTRACT-CYCLES₃ from Theorem 4.1 is inherited by 2-EXCHANGE.

For the lower bound on the performance guarantee, given the graph in Fig. 2, 2-EXCHANGE can choose a number of edges arbitrarily close to 1.75 times the minimum. There are $\frac{n}{4}$ groups with four vertices in each group. First observe that the graph has a directed Hamilton cycle. The edges marked in Fig. 2 form a solution with which 2-EXCHANGE could terminate. This solution clearly has $\frac{7n}{4}$ edges. This gives the lower bound of 1.75 on the performance of the algorithm. □

6. Implementation. For any fixed k , CONTRACT-CYCLES _{k} can be implemented in polynomial time using exhaustive search to find long cycles. For instance, if a cycle of size at least k exists, one can be found in polynomial time as follows: For each simple path P of $k - 1$ edges, check whether a path from the head of P to the tail exists after P 's internal vertices are removed from the graph. If k is even, there are at most $m^{k/2}$ such paths; if k is odd, the number is at most $n m^{(k-1)/2}$. It takes $O(m)$ time to decide if there is a path from the head of P to the tail of P . For the first iteration of the for loop, we may have $O(n)$ iterations of the while loop. Since the first iteration is the most time consuming, the algorithm can be implemented in $O(n m^{1+k/2})$ time for even k and $O(n^2 m^{(k+1)/2})$ time for odd k .

6.1. A practical implementation yielding 1.75. Next we give a practical, near linear-time implementation of CONTRACT-CYCLES₃. The performance guarantee achieved is $c_3 = 1.75$. CONTRACT-CYCLES₃ consists of two phases: (1) repeatedly finding and contracting

cycles of three or more edges (called *long cycles*) until no such cycles exist, and (2) contracting the remaining 2-cycles.

High-level description of the algorithm. To perform phase (1), the algorithm does a DFS of the graph from an arbitrary root. During the search, the algorithm identifies edges for contraction by adding them to a set S . At any point in the search, G' denotes the subgraph of edges and vertices traversed so far. The rule for adding edges to S is as follows: when a new edge is traversed, if the new edge creates a long cycle in G'/S , the algorithm adds the edges of the cycle to S . The algorithm thus maintains that G'/S has no long cycles. When the DFS finishes, G'/S has only 2-cycles. The edges on these 2-cycles, together with S , are the desired SCSS.

Because G'/S has no long cycles and the original graph is strongly connected, G'/S maintains a simple structure.

LEMMA 6.1. *After the addition of any edge to G' and the possible contraction of a cycle by adding it to S , (i) the graph G'/S consists of an outward branching and some of its reverse edges, (ii) the only reverse edges that might not be present are those on the “active” path: from the supervertex containing the root to the supervertex in G'/S containing the current vertex of the DFS.*

Proof. Clearly the invariant is initially true. We show that each given step of the algorithm maintains the invariant. In each case, if u and w denote vertices in the graph, then let U and W denote the vertices in G'/S containing u and w , respectively.

When the DFS traverses an edge (u, w) to visit a new vertex w , we have the following: Vertex w and edge (u, w) are added to G' . Vertex w becomes the current vertex. In G'/S , the outward branching is extended to the new vertex W by the addition of edge (U, W) . No other edge is added and no cycle is created. Thus, part (i) of the invariant is maintained. The supervertex containing the current vertex is now W , and the new “active path” contains the old “active path.” Thus, part (ii) of the invariant is also maintained.

When the DFS traverses an edge (u, w) and w is already visited, we have the following: If $U = W$ or the edge (U, W) already exists in G'/S , then no cycle is created, G'/S is unchanged, and the invariant is clearly maintained. Otherwise, the edge (u, w) is added to G' and a cycle with the simple structure illustrated in Fig. 3 is created in G'/S . The cycle consists of the edge (U, W) followed by the (possibly empty) path of reverse edges from W to the lowest common ancestor (lca) of U and W , followed by the (possibly empty) path of branching edges from $\text{lca}(U, W)$ to U . Addition of (U, W) to G'/S and contraction of this cycle (in case it is a long cycle) maintains part (i) of the invariant. If the “active path” is changed, it is only because part of it is contracted, so part (ii) of the invariant is maintained.

When the DFS finishes visiting a vertex w , we have the following: No edge is added and no cycle is contracted, so part (i) is clearly maintained. Let u be the new current vertex, i.e., w 's parent in the DFS tree. If $U = W$, then part (ii) is clearly maintained. Otherwise, consider the set D of descendants of w in the DFS tree. Since the original graph is strongly connected, some edge (x, y) in the original graph goes from the set D to its complement $V - D$. All vertices in D have been visited, so (x, y) is in G' . By part (i) of the invariant, the vertex in G'/S containing x must be W , while the vertex in G'/S containing y must be U . Otherwise the edge corresponding to (x, y) in G'/S would create a long cycle. \square

The algorithm maintains the contracted graph G'/S , using a union-find data structure [22] to represent the vertices in the standard way and three data structures to maintain the branching, the reverse edges discovered so far, and the “active path.” When a cycle arises in G'/S , it must be of the form described in the proof of Lemma 6.1 and illustrated in Fig. 3. Using these data structures, the algorithm discovers it and, if it is long, contracts it in a number of union-find operations proportional to the length of the cycle. This yields an $O(m\alpha(m, n))$ -time algorithm.




FIG. 3. Contracted graph G'/S .

The vertices of G'/S are represented in union-find sets as follows:

MAKE-SET(v). Adds the set $\{v\}$ corresponding to the new vertex of G'/S .

FIND(v). Returns the set in G'/S that contains vertex v .

UNION(u, v). Joins into a single set the two sets corresponding to the vertices in G'/S containing G' 's vertices u and v .

The data structures representing the branching, reverse edges, and the active paths, respectively are as follows:

from-root[W]. For each branching edge (U, W) in G'/S , from-root[W] = (u, w) for some $(u, w) \in (U \times W) \cap E$.

to-root[U]. For each reverse edge (U, W) in G'/S , to-root[U] = (u, w) for some $(u, w) \in (U \times W) \cap E$.

to-active[U]. For each vertex U on the "active path" in G'/S , to-active[U] = (u, w) , where $(u, w) \in (U \times W) \cap E$ and W is the child of U for which the recursive DFS call is currently executing, unless no recursive DFS is executing, in which case to-active[U] = **current**.

For all other vertices, to-active[U] = **nil**.

Pseudo code for the algorithm is given in Figs. 4 and 5.

By the preceding discussion, the algorithm implements CONTRACT-CYCLES₃. It is straightforward to show that it runs in $O(m\alpha(m, n))$ time. Hence, we have the following theorem.

THEOREM 6.2. *There is an $O(m\alpha(m, n))$ -time approximation algorithm for the minimum SCSS problem achieving a performance guarantee of 1.75 on an m -edge, n -vertex graph.*

Here $\alpha(m, n)$ is the inverse-Ackermann function associated with the union-find data structure [22].

Example to illustrate algorithm. In the example in Fig. 6, the algorithm begins the DFS from vertex 1. It visits vertices 2,3,4 and then traverses the reverse edge (4, 2). Since this edge creates a 3-cycle (2, 3), (3, 4), (4, 2) in G'/S , it contracts the cycle. Next it traverses the reverse edge (3, 1), but does not contract it, since it forms only a 2-cycle in the contracted graph. Continuing the DFS, it visits vertices 5 and 6. When it traverses the edge (6, 4), it discovers and contracts the cycle (3, 1), (1, 5), (5, 6), (6, 4). Next it visits vertices 7 and 8, traversing the reverse edges (8, 7) and (7, 6). Traversing the edge (6, 8), it discovers and

CONTRACT-CYCLES₃($G = (V, E)$) — *Pseudocode.*

```

1   $S \leftarrow \{\}$ 
2  Choose  $r \in V$ .
3  DFS( $r$ )
4  Add 2-cycles remaining in  $G'/S$  to  $S$ .
5  return  $S$ 

DFS( $u$ ) —
1  to-active[FIND( $u$ )]  $\leftarrow$  current
2  for each vertex  $w$  adjacent to  $u$  — traverse edge ( $u, w$ ) —
3    if ( $w$  is not yet visited) — new vertex —
4      MAKE-SET( $w$ )
5      to-active[FIND( $u$ )]  $\leftarrow$  from-root[FIND( $w$ )]  $\leftarrow$  ( $u, w$ )
6      DFS( $w$ )
7      to-active[FIND( $u$ )]  $\leftarrow$  current
8    else — edge creates cycle in  $G'/S$  —
9      if (FIND( $u$ )  $\neq$  FIND( $w$ )) — cycle length at least 2 —
10       ( $x, y$ )  $\leftarrow$  from-root[FIND( $u$ )]
11       if (FIND( $x$ ) = FIND( $w$ )) — length two cycle through parent,  $U - W - U$  —
12         to-root[FIND( $u$ )]  $\leftarrow$  ( $u, w$ ) — record edge to parent —
13       else
14         ( $x, y$ )  $\leftarrow$  from-root[FIND( $w$ )]
15         if (FIND( $x$ )  $\neq$  FIND( $u$ )) — not a forward edge to child; length of cycle  $\geq 3$  —
16           CONTRACT-CYCLE( $w$ )
17            $S \leftarrow S \cup \{(u, w)\}$ 
18 to-active[FIND( $u$ )]  $\leftarrow$  nil

```

FIG. 4. Practical implementation of CONTRACT-CYCLES₃.

CONTRACT-CYCLE(w) —

```

1  while (to-active[FIND( $w$ )]  $\neq$  current) do
2    if (to-active[FIND( $w$ )] = nil) then — Go up towards lca along reverse edges. —
3      ( $c, p$ )  $\leftarrow$  to-root[FIND( $w$ )]
4       $a \leftarrow$  to-active[FIND( $p$ )]
5    else — Go down from lca along active path. —
6      ( $p, c$ )  $\leftarrow$  to-active[FIND( $w$ )]
7       $a \leftarrow$  to-active[FIND( $c$ )]
      — Contract parent  $p$  and child  $c$ . —
8       $f \leftarrow$  from-root[FIND( $p$ )]
9       $t \leftarrow$  to-root[FIND( $p$ )]
10     UNION( $p, c$ )
11     to-active[FIND( $w$ )]  $\leftarrow$   $a$ 
12     from-root[FIND( $w$ )]  $\leftarrow$   $f$ 
13     to-root[FIND( $w$ )]  $\leftarrow$   $t$ 

```

FIG. 5. Subroutine CONTRACT-CYCLE.

contracts the 3-cycle (8, 7), (7, 6), (6, 8). In this example, no 2-cycles remain, so it returns just the contracted edges.

7. Potential improvement of CONTRACT-CYCLES_k. A natural modification to CONTRACT-CYCLES_k would be to stop when the contracted graph has no cycles of length more than some c and somehow solve the remaining problem optimally.

For instance, for $c = 3$, by following the proof of Theorem 4.1, one can show that this would improve the performance guarantee of CONTRACT-CYCLES_k to $c_k - \frac{1}{36}$ (for $k \geq 4$), matching the lower bound in Table 1. (The lower bound given holds for the modified algorithm.)




FIG. 6. Example for illustrating execution of algorithm.

This leads us to consider the minimum $SCSS_c$ problem—the minimum SCSS problem restricted to graphs with cycle length bounded by c . The following theorem is shown in [14].

THEOREM 7.1. *There is a polynomial-time algorithm for the $SCSS_3$ problem.*

We make no conjecture concerning the $SCSS_4$ problem. However, we next show that the $SCSS_5$ problem is NP-hard and the $SCSS_{17}$ problem is MAX SNP-hard.

7.1. NP-hardness of $SCSS_5$. We prove the following theorem.

THEOREM 7.2. *The minimum $SCSS_5$ problem is NP-hard.*

Proof. The proof is by a reduction from SAT [8]. Fix an arbitrary formula F in conjunctive normal form (CNF). We will build a rooted digraph such that any SCSS contains all the edges out of the root (d of them) and F is satisfiable if and only if there exists an SCSS E' in which each of the remaining $n - 1$ nonroot vertices has out-degree equal to one. Thus the formula will be satisfiable if and only if there is an SCSS with $n - 1 + d$ edges.

The graph has a fixed root vertex r and a vertex for each clause in F (these vertices are not shown in Fig. 7). Each clause vertex has a return edge to the root. For each variable in F , the graph has an instance of the gadget illustrated in Fig. 7. The edges into the gadget come from the root. Each such edge is present in any SCSS. The edges out of the gadget are alternately labeled $+$ and $-$. For every clause with a positive instance of the variable, one of the $+$ edges goes to the clause vertex. For every clause with a negative instance of the variable, one of the $-$ edges goes to the clause vertex. Unassigned $+$ and $-$ edges go to the root. (The gadget is easily enlarged to allow any number of occurrences.)




FIG. 7. Variable gadget for NP-hardness proof.

The key property of the gadget is that if every (nonroot) vertex has out-degree one in some SCSS, then either all of the counterclockwise edges are in the SCSS (corresponding to the variable being true) or all of the clockwise edges are in the SCSS (corresponding to the variable being false). Thus, given any SCSS of $d + n - 1$ edges, where d is the out-degree of the root and n is the number of vertices in the digraph constructed, it is easy to construct a satisfying assignment for F . Conversely, given any satisfying assignment for F , it is easy to construct an SCSS of size $d + n - 1$. \square

7.2. MAX SNP-hardness of SCSS₁₇. Next we consider the MAX SNP-hardness of the problem. To prove this we do a reduction from the vertex-cover problem in bounded-degree graphs to the SCSS problem. Since the proof closely follows the reduction from vertex cover to Hamiltonian circuits (see [8]), it is suggested that the reader study this reduction before reading this section. It is known that the problem of finding a minimum vertex cover is MAX SNP-hard in graphs whose maximum degree is bounded by seven [19].

Let G be a connected, undirected graph whose maximum degree is bounded by seven. Let G have m edges and n vertices. We construct a digraph D with $2m + 1$ vertices and no cycle longer than 17. Any vertex cover of G of size s will yield an SCSS in D of size $2m + s$ and vice versa. We then show that, since G has $O(n)$ edges, this yields an L-reduction (i.e., an approximation-preserving reduction [19]).

7.2.1. The construction of D . Applying Vizing’s theorem [23], color the edges of G in polynomial time with at most eight colors so that no two edges incident to a vertex share the same color. Let the colors of the edges be chosen from the set $\{1, 2, \dots, 8\}$.

The construction begins with a special “root” vertex r in D .

As the construction proceeds, each vertex in G will have a “current vertex” in D , initially the root vertex. We process edges in each color class starting with color 1. For each edge (u, v) , add a “cover-testing gadget” to D as illustrated in Fig. 8. Specifically, add two new vertices x and y . Add two edges into x : the first, labeled u^+ , from the current vertex of u ; the




FIG. 8. A cover-testing component.

second, labeled u^- , from y . Similarly, add two edges into y : the first, labeled v^+ , from the current vertex of v ; the second, labeled v^- , from x . Make y the new current vertex of u ; make x the new current vertex of v . Finally, after all edges of G have been considered, for each vertex v in G , add an edge labeled v^+ from its final current vertex to the root. The gadgets are implicitly layered with each gadget assigned to a layer corresponding to the color of the associated edge in G .

LEMMA 7.3. *The graph D constructed above has no cycle with more than 17 edges.*

Proof. We first assign numbers to the vertices of D . The root r is assigned the number 0. The construction above proceeds in order of increasing color of the edges of G . When considering an edge (u, v) of color c , we add two new vertices: x is added to v 's path and y is added to u 's path. We assign the vertices x and y the number c . Consider any cycle X of length greater than two in D . It is clear that such a cycle must pass through r , since D is layered. Hence the cycle is of the form $(r, x_1, x_2, \dots, x_k, r)$. Because we considered the edges in order of increasing color, the numbers assigned to the vertices in X increase at least every two steps in any path in D (not including r). In other words, the numbers assigned to the vertices x_1, \dots, x_k form a nondecreasing sequence in which no three consecutive vertices get the same number. Since the edges of G were colored with 8 colors, the numbers assigned to the vertices of D range from 0 to 8 (only r gets the number 0). Combining all these, the length of the cycle X is at most 17. \square

7.2.2. The analysis. We now show that every vertex cover of G has a corresponding SCSS in D . The proof is similar to the corresponding proof (in the reduction from vertex cover to Hamiltonian circuits) that every vertex cover has a corresponding Hamiltonian circuit. Consider an arbitrary vertex cover S of G . The idea is to choose the paths in the SCSS corresponding to S in D . The paths of the vertices of $V - S$ are not yet connected. Since S forms a vertex cover, the vertices in the paths of $V - S$ can be connected using the cover-testing components.

LEMMA 7.4. *Given a vertex cover of size s in G , an SCSS of D of size $2m + s$ can be constructed.*

Proof. Construct a subgraph H of D as follows. For each vertex u in G , let d_u be the degree of u in G . If u is in the vertex cover, add the $d_u + 1$ edges labeled u^+ in D to H . Otherwise, add the d_u edges labeled u^- in D to H . It is easy to verify that H has the following properties:

1. H has $2m + s$ edges.
2. H has no cycles of length 2.
3. Every vertex of H has at least one outgoing and one incoming edge.

As mentioned earlier, D is layered and every cycle of length greater than 2 contains r . Therefore property 2 above implies that every cycle of H passes through r . By the above conditions, H

contains a path from r to every vertex v and another path from v to r , and is therefore strongly connected. To obtain a path from any v to r , start from v and keep traversing an outgoing edge (which exists by property 3) from the current vertex. Such a path must eventually reach r because r is contained in every cycle of H . Hence H satisfies the lemma. \square

We now show that every SCSS of D corresponds to a vertex cover of G . The proof works by showing that any SCSS can be converted into a “canonical” SCSS, whose size is no larger, that corresponds to a vertex cover of G .

LEMMA 7.5. *Given an SCSS in D of size $2m + s$, a vertex cover of G of size s can be constructed.*

Proof. As long as some nonroot vertex y has both of its incoming edges in the SCSS, modify the SCSS as follows: Let (x, y) be the edge labeled v^- for some v . Remove the edge (x, y) and add the other edge out of x if it is not already present. Alternatively, if some nonroot vertex x has both of its outgoing edges in the SCSS, remove the edge (x, y) and add the other edge into y . Repeat either modification as long as applicable.

By the layering of D , each modification maintains the strong connectivity of the SCSS. Clearly none of the modifications increase the size. Each step reduces the number of edges labeled u^- for some u in the SCSS, so after at most $2m$ steps neither modification applies, and in the resulting SCSS every nonroot vertex has exactly one incoming edge and one outgoing edge in the SCSS.

An easy induction on the layering shows that for any vertex v in G , either all of the edges labeled v^+ in D are in the SCSS or none are in the SCSS, in which case all of the edges labeled v^- are in the SCSS. Let C be the set of vertices in G of the former kind. It is easy to show that the size of the SCSS is $2m + |C|$, so that $|C| \leq s$. For every edge (u, v) in G , the form of the gadget ensures that at least one of the two endpoints is in C . Hence, C is the desired cover. \square

THEOREM 7.6. *The minimum SCSS₁₇ problem is MAX SNP-hard.*

Proof. Let G be an arbitrary undirected graph G whose maximum degree is bounded by seven. Let G have m edges and n vertices. Construct the digraph D as shown earlier. By Lemma 7.3, D has no cycles greater than 17. By Lemma 7.4, any vertex cover in G of size s can be used to obtain an SCSS of D of size $2m + s$. Conversely, by Lemma 7.5 an SCSS of D of size $2m + s$ can be used to obtain a vertex cover of G of size s . Since the degree of G is bounded, $m = O(n) = O(s)$ and it is easily verified that this yields an L-reduction from degree-bounded vertex cover to the minimum SCSS₁₇ problem. \square

8. Open problems. An obvious problem is to further characterize the various complexities of the minimum SCSS _{k} problems.

The most interesting open problem is to obtain a performance guarantee that is less than 2 for the weighted strong connectivity problem (as mentioned earlier, the performance factor of 2 is from Frederickson and Jájá [6]). Such an algorithm may have implications for the weighted 2-connectivity problem [15] in undirected graphs as well.

The performance guarantee of k -EXCHANGE probably improves as k increases. Proving this would be interesting—similar “local improvement” algorithms are applicable to a wide variety of problems.

Acknowledgments. We thank the referees for useful comments.

REFERENCES

- [1] A. AGRAWAL, P. KLEIN, AND R. RAVI, *When trees collide: An approximation algorithm for the generalized Steiner problem on networks*, Proc. 23rd ACM Symposium on Theory of Computing, New Orleans, LA, 1991, pp. 134–144.

- [2] A. V. AHO, M. R. GAREY, AND J. D. ULLMAN, *The transitive reduction of a directed graph*, SIAM J. Comput., 1 (1970), pp. 131–137.
- [3] S. ARNBORG, J. LAGERGREN, AND D. SEESE, *Easy problems for tree-decomposable graphs*, J. Algorithms, 12 (1991), pp. 308–340.
- [4] S. ARORA, C. LUND, R. MOTWANI, M. SUDAN, AND M. SZEGEDY, *Proof verification and hardness of approximation problems*, Proc. 33rd IEEE Symposium on Foundations of Computer Science, Pittsburgh, PA, 1992, pp. 14–23.
- [5] T. H. CORMEN, C. E. LEISERSON, AND R. L. RIVEST, *Introduction to Algorithms*, MIT Press, Cambridge, MA, 1989.
- [6] G. N. FREDERICKSON AND J. JÁJÁ, *Approximation algorithms for several graph augmentation problems*, SIAM J. Comput., 10 (1981), pp. 270–283.
- [7] H. N. GABOW, Z. GALIL, T. SPENCER, AND R. E. TARJAN, *Efficient algorithms for finding minimum spanning trees in undirected and directed graphs*, Combinatorica, 6 (1986), pp. 109–122.
- [8] M. R. GAREY AND D. S. JOHNSON, *Computers and intractability: A guide to the theory of NP-completeness*, W. H. Freeman, San Francisco, CA, 1979.
- [9] N. GARG, V. SANTOSH, AND A. SINGLA, *Improved approximation algorithms for biconnected subgraphs via better lower bounding techniques*, Proc. 4th Annual ACM–SIAM Symposium on Discrete Algorithms, Austin, TX, 1993, pp. 103–111.
- [10] P. GIBBONS, R. M. KARP, V. RAMACHANDRAN, D. SOROKER, AND R. E. TARJAN, *Transitive compaction in parallel via branchings*, J. Algorithms, 12 (1991), pp. 110–125.
- [11] M. GOEMANS AND D. WILLIAMSON, *A general approximation technique for constrained forest problems*, Proc. 3rd Annual ACM–SIAM Symposium on Discrete Algorithms, Orlando, FL, 1992, pp. 307–316.
- [12] F. HARARY, R. Z. NORMAN, AND D. CARTWRIGHT, *Structural models: An introduction to the theory of directed graphs*, John Wiley, New York, NY, 1965.
- [13] H. T. HSU, *An algorithm for finding a minimal equivalent graph of a digraph*, J. Assoc. Comput. Mach., 22 (1975), pp. 11–16.
- [14] S. KHULLER, B. RAGHAVACHARI, AND N. YOUNG, *On strongly connected digraphs with bounded cycle length*, UMIACS-TR-94-10/CS-TR-3212, University of Maryland, 1994.
- [15] S. KHULLER AND U. VISHKIN, *Biconnectivity approximations and graph carvings*, Proc. 24th ACM Symposium on Theory of Computing, 1992, pp. 759–770; J. Assoc. Comput. Mach., 41 (1994), pp. 214–235.
- [16] P. N. KLEIN AND R. RAVI, *When cycles collapse: A general approximation technique for constrained two-connectivity problems*, Proc. 3rd Integer Programming and Combinatorial Optimization Conference, Mathematical Programming Society, Erice, Italy, 1993, pp. 39–56.
- [17] D. E. KNUTH, *Fundamental Algorithms*, Addison–Wesley, Menlo Park, CA, 1973.
- [18] D. M. MOYLES AND G. L. THOMPSON, *An algorithm for finding the minimum equivalent graph of a digraph*, J. Assoc. Comput. Mach., 16 (1969), pp. 455–460.
- [19] C. H. PAPADIMITRIOU AND M. YANNAKAKIS, *Optimization, approximation, and complexity classes*, J. Comput. System Sci., 43 (1991), pp. 425–440.
- [20] S. SAHNI, *Computationally related problems*, SIAM J. Comput., 3 (1974), pp. 262–279.
- [21] K. SIMON, *Finding a minimal transitive reduction in a strongly connected digraph within linear time*, in Proc. 15th International Workshop WG'89, Lecture Notes Comput. Sci. 411, Springer-Verlag, 1989, pp. 245–259.
- [22] R. E. TARJAN, *Data Structures and Network Algorithms*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 1983.
- [23] V. G. VIZING, *On an estimate of the chromatic class of a P-graph*, Diskret. Anal., 3 (1964), pp. 25–30. (In Russian.)
- [24] D. P. WILLIAMSON, M. X. GOEMANS, M. MIHAIL, AND V. V. VAZIRANI, *A primal-dual approximation algorithm for generalized Steiner network problems*, Proc. 25th ACM Symposium on Theory of Computing, San Diego, CA, 1993, pp. 708–717.