Wavelength Conversion in Optical Networks*

Jon Kleinberg [†] Amit Kumar [‡]

Abstract

In many models of optical routing, we are given a set of *communication paths* in a network, and we must assign a wavelength to each path so that paths sharing an edge receive different wavelengths. The goal is to assign as few wavelengths as possible, in order to make as efficient use as possible of the optical bandwidth. Much work in the area of optical networks has considered the use of *wavelength converters*: if a node of a network contains a wavelength converter, any path that passes through this node may change its wavelength. Having converters at some of the nodes can reduce the number of wavelengths required for routing, down to the following natural *congestion bound*: even with converters, we will always need at least as many wavelengths as the maximum number of paths sharing a single edge. Thus Wilfong and Winkler defined a set S of nodes in a network to be *sufficient* if, placing converters at the nodes in S, every set of paths can be routed with a number of wavelengths equal to its congestion bound. They showed that finding a sufficient set of minimum size is NP-complete, even in planar graphs.

In this paper, we provide a polynomial-time algorithm to find a sufficient set for an arbitrary directed network whose size is within a factor of 2 of minimum. We also observe that improving on the factor of 2 would lead to a corresponding improvement for the vertex cover problem. For the case of planar graphs, we provide a polynomial-time approximation scheme. The algorithms are based on connections between the minimum sufficient set problem and the *undirected feedback vertex set problem*. In particular, as a component of the algorithm on planar graphs, we develop the first polynomial-time approximation scheme for the undirected feedback vertex set problem in planar graphs, a result that we feel to be of interest in its own right.

1 Introduction

The assignment of wavelengths to communication paths is a basic optimization problem for optical networks based on *wavelength division multiplexing* (WDM) [12, 13, 22]. The problem can be modeled in the following standard way (e.g. [1, 21, 24]). We are given a directed graph G, and a set $\mathcal{P}$ of paths in G, and wish to assign a *color* (or *wavelength*) to each path in $\mathcal{P}$, so that no two paths sharing an edge receive the same color; we will call such a wavelength assignment *valid*. The goal is to minimize the number of colors used in a valid assignment; we will denote this minimum by $\chi(\mathcal{P})$.

---

*A preliminary version of this paper appears in the Proceedings of the 10th ACM-SIAM Symposium on Discrete Algorithms, 1999.

[†]Department of Computer Science, Cornell University, Ithaca NY 14853. Email: kleinber@cs.cornell.edu.
Supported in part by an Alfred P. Sloan Research Fellowship and by NSF Faculty Early Career Development Award CCR-9701399.

[‡]Department of Computer Science, Cornell University, Ithaca NY 14853. Email: amitk@cs.cornell.edu.

Much work in the area of optical networks has considered the use of *wavelength converters*, as a way to make more efficient use of optical bandwidth [7, 12, 15, 17, 21, 25]. Placing a wavelength converter at a node v enables any path containing v to change its color as it passes through v. In a network with converters, our notion of a valid wavelength assignment must become more general: it is now an assignment of a wavelength to each *edge* of each path, with the restriction that the sequence of color assignments to a path can only change when the path passes through a converter. Note that if we were to place a converter at every node of G, the minimum number of colors required in a valid assignment would be equal to the following natural *congestion bound* $\nu(\mathcal{P})$: the maximum number of paths passing through any single edge. In the absence of converters, $\chi(\mathcal{P})$ can be arbitrarily larger than $\nu(\mathcal{P})$ on some instances; even with converters, we cannot get away with fewer than $\nu(\mathcal{P})$ colors, as a different color is required for each path that passes through the most congested edge.

Since wavelength converters are expensive components of a network, one would like to use them as parsimoniously as possible while still achieving a substantial reduction in wavelength usage. Motivated by this, Wilfong and Winkler [27] defined a subset S of the nodes of a graph to be *sufficient* if it has the following property: with converters at the nodes in S, any set $\mathcal{P}$ of paths has a valid wavelength assignment with only $\nu(\mathcal{P})$ colors. They then proposed the following basic network design problem: Given a graph G, find a sufficient set of minimum size. They proved this problem to be NP-complete, even on planar graphs [27]; the existence of good approximation algorithms was left open.

Approximating the Minimum Sufficient Set. Our first main result is a polynomial-time 2-approximation for the minimum sufficient set, in an arbitrary directed graph. Given the directed nature of the problem, and the (easily proven) fact that one requires at least one converter on every directed cycle of G, it may initially appear that the problem is closely related to the *directed feedback vertex set problem*, for which constant-factor approximation algorithms are not known. However, we show how to solve it using techniques based on the *undirected* feedback set problem [4, 8, 9]. (Recall that a feedback set in an undirected (resp. directed) graph is a subset of the vertices that intersects all undirected (resp. directed) cycles.) Our techniques shed further light on the nature of obstructions to the efficient use of optical bandwidth in routing problems, through the analysis of configurations that necessitate the use of wavelength converters.

We also show by a simple construction that any polynomial-time c-approximation for the minimum sufficient set, with $c < 2$, would imply the existence of a polynomial-time c-approximation for the minimum vertex cover in an undirected graph, which remains a long-standing open question.

Sufficient Sets and Feedback Sets in Planar Graphs. Our second main result is a polynomial-time approximation scheme for the minimum sufficient set in an arbitrary directed planar graph. Building on the techniques developed for the 2-approximation algorithm, we obtain this result by designing a polynomial-time approximation scheme for the feedback vertex set problem in undirected planar graphs.

We feel that this latter result is of interest in its own right, as it resolves a question that has remained open despite several recent investigations of the undirected feedback set problem [4, 6, 8, 9, 11]. In particular, Bar-Yehuda et al. [6] obtained a 10-approximation for the weighted version of the problem in planar graphs, and Goemans and Williamson [11] showed that the natural *cycle formulation* of the problem as an integer program has an integrality gap of at most $\frac{9}{4}$. However, no result prior to ours for planar graphs improved on the approximation ratio of 2, which holds for general graphs.

The difficulty in designing an approximation scheme for the undirected planar feedback set problem stems in part from the fact that neither the Lipton-Tarjan approach via planar separators [18, 19] nor Baker's approach via outerplanar decompositions [5] is directly applicable. Specifically, the Lipton-Tarjan method incurs an $O(n/\sqrt{\log n})$ additive error, so one needs to be able to guarantee that the optimum is of near-linear size before invoking this approach. Baker's method works on a large number of independent "strips" in the graph, each of bounded tree-width; in our case, a single cycle may span many of these strips and hence not be considered in any of the independent sub-problems.

Our approach is to invoke a Lipton-Tarjan-based method if the graph G has bounded degree. In particular, after deleting bridge edges and "dissolving" degree-2 vertices, one can show via an analysis of the intersection graph of the faces that the minimum feedback set must have linear size. If G has high-degree vertices, one cannot necessarily include all of them in constructing a near-optimal solution. However, we show that including all high-degree vertices is possible provided that G does not contain a certain type of "bad configuration"; and if there is such a configuration, it provides us with a way to decompose G so as to continue by recursion.

Other Related Work. In addition to the references discussed above on wavelength converters, sufficient sets, and feedback sets, there has been a line of research directed at developing approximation algorithms for routing problems in optical networks; see e.g. [3, 14, 16, 20, 23, 24]. The setting of these papers differs from what we consider here in several respects. They are concerned with the problem of choosing paths and assigning wavelengths for a collection of terminal pairs in a network; we focus on a network design question, ensuring a property that must hold for all possible sets of terminal pairs. Note also that our notion of wavelength assignment (following the model of Wilfong and Winkler [27]) is based on a set of terminal pairs *with specified paths*; much of the previous algorithmic work on these questions has considered a model that couples path selection and wavelength assignment.

There has also been work on the problem of routing and wavelength assignment in networks with a given subset of the nodes designated as converters; see for example [2, 25] and the references therein. One issue that they study, which we do not address here, is the capabilities of converters that are only able to convert a fixed, bounded number of wavelengths.

2 Preliminaries

A *network*, in this paper, is a directed graph $G = (V, E)$. For certain pairs of nodes u and v, G may contain edges (u, v) and (v, u); this will be called a *bi-directed pair* of edges. G will be called *bi-directed* if each edge belongs to a bi-directed pair. The *skeleton* of the network G, denoted $s(G)$, is the undirected graph obtained from G by ignoring the directions of the edges, and replacing each bi-directed pair of edges by a single undirected edge. For the sake of simplicity, we will assume that $s(G)$ is connected for all networks G that we deal with. We will also assume that G does not contain multiple copies of the same edge in the same direction.

Let $\mathcal{P}$ denote a set of directed paths in G. We will assume throughout that no individual path in $\mathcal{P}$ passes through any vertex more than once. The *congestion* of $\mathcal{P}$, denoted $\nu(\mathcal{P})$, is defined to be the maximum, over directed edges e of G, of the number of paths containing e. The *conflict graph* of $\mathcal{P}$ is an undirected graph K whose vertices are the set of paths in $\mathcal{P}$; two paths P_i and P_j are joined by an edge in K if and only if they share an edge in G. $\chi(\mathcal{P})$ — the minimum number of colors in a valid wavelength assignment for $\mathcal{P}$ — is simply the chromatic number of the conflict graph of $\mathcal{P}$.

3 Sufficient Sets in General Graphs

3.1 Bi-Directed Graphs

We begin with a brief description of the 2-approximation for bi-directed graphs, since it is conceptually simple, and it represents the case considered by Wilfong and Winkler [27].

Define a vertex v to be a *branching node* if its degree in $s(G)$ is greater than 2. Analogously, define a *straight node* to be a node whose degree in $s(G)$ is less than or equal to 2. (We will also refer to a node of degree 1 as a *leaf.*) We will assume that $s(G)$ contains at least one branching node, since otherwise $s(G)$ is either a path or a cycle, and Minimum Sufficient Set is easy to solve optimally. We say that a node of a path P is an *internal node* if it is not one of the two endpoints.

Wilfong and Winkler [27] developed some basic properties of converters in a bi-directed graph; we restate one of their results here, as it will be useful in what follows. We begin with some definitions from [27]. For a subset S of V, define $G_s(S) = (V(S), E'(S))$ as the graph in which $V(S)$ is the set of nodes in $V - S$, together with the set of pairs (s, e) for which $s \in S$ and $e \in E'$ is an edge incident on s. $E'(S)$ consists of (i) edges of G_s between vertices in $V - S$; (ii) edges of the form $((s, e), v)$ whenever $(s, e) \in V(S)$ and $e = (s, v)$; and (iii) $((s, e), (t, e))$ where s and t are in S. Define a *spider* to be a tree with at most one vertex of degree greater than 2. Now the following theorem is proved in [27].

Theorem 3.1 *Let G be a bi-directed graph. A set S of nodes is sufficient for G if and only if each component of $G(S)$ is a spider.*

We will say that a sufficient set is *canonical* if it contains only branching nodes.

Proposition 3.2 *Let G be a bi-directed graph. If S is a sufficient set for G, then there exists another sufficient set S' for G such that $|S'| \leq |S|$ and S' (as a subset of vertices in G_s) doesn't contain any straight node.*

We construct a graph $H = (V_H, E_H)$ as follows. V_H consists of all branching nodes in $s(G)$. For $u, v, \in V_H$, (u, v) is an edge in E_H if and only if there exists a path in $s(G)$ between

u and v such that all internal nodes in the path are straight nodes. Note that H may have self-loops, which we retain as part of the graph.

Theorem 3.3 *Let S be a canonical sufficient set, and consider S as a subset of V_H. Then S is a vertex cover of H. Conversely, every vertex cover of H is also a sufficient set of G. (By definition, we require a vertex cover to include the unique endpoint of each self-loop.)*

Proof. Suppose S is a canonical sufficient set of G, but is not a vertex cover of H. Then there exists an edge (u, v) in H such that neither u nor v are in S. Now, in $s(G)$, u and v are branching nodes such that there is a path P between them containing no other branching node. Since S cannot contain any straight node, none of the nodes in P connecting u and v are in S. When we form the graph $G_s(S)$, all the nodes in P — including u and v — will be in the same component. Moreover, the degrees of both u and v will be the same as that in G. Thus, if $u = v$, this would be a cycle in a component of $G_s(S)$; if $u \neq v$, this would be two branching nodes in a component of $G_s(S)$. In both cases, this is a contradiction. So, S must be a vertex cover of H.

Conversely, suppose C is a vertex cover for H, but not a sufficient set of G. Then there is a component Z of $G_s(C)$ that is not a spider. Note that any non-leaf node of Z corresponds to a node of the same degree in $V - C$. If Z contains a single branching node v, then Z must contain a cycle including v. But this implies that v is the endpoint of a self-loop of H; since we also know $v \notin C$, this contradicts the assumption that C is a vertex cover. Otherwise, Z contains more than one branching node. Choose a pair $u, v \in Z$ of branching nodes whose distance in Z is minimal. Thus, there is a u-v path P in Z whose internal nodes all have degree 2. These nodes must also have degree two in $s(G)$, and hence (u, v) is an edge of H. But this again contradicts the assumption that C is a vertex cover of H. ■

By applying Theorem 3.3 and a 2-approximation for the vertex cover problem, we have

Corollary 3.4 *There is a polynomial-time* 2*-approximation for the minimum sufficient set problem in bi-directed graphs.*

One can also use the approximation scheme for vertex cover ([5],[19]) in planar graphs to get an approximation scheme for the minimum sufficient set problem in bi-directed planar graphs.

A simple adaptation of the proof of Theorem 3.3 also shows the following.

Proposition 3.5 *If there is a polynomial-time c-approximation for the minimum sufficient set problem, then there is a polynomial-time c-approximation for the minimum vertex cover problem.*

3.2 General Graphs

The case of general directed graphs is significantly more complicated, and it requires more work to obtain a characterization of the configurations that require a wavelength converter.

First note that on a directed cycle, one can place a set $\mathcal{P}$ of three paths so that $\nu(\mathcal{P}) = 2$, but each pair of paths meets at some edge, and hence $\chi(\mathcal{P}) = 3$. Thus, any sufficient set must contain at least one node from each directed cycle of G. We now look at a different type of subgraph that also necessitates a converter. Consider a graph consisting of a bi-directional path with edges coming in and out at its two ends. For concreteness, let the path be P and its ends be u and v. Suppose that we have edges coming in and out at both ends, i.e., (u, u_1) and (u_2, u) are edges incident on u ($u_1 \neq u_2$). Similarly, $(v, v_1), (v_2, v)$ are incident on v ($v_1 \neq v_2$). Let us call such a graph an *$\mathcal{H}$-graph.* The bi-directional path P of an $\mathcal{H}$-graph will be called its *characteristic path.*

As shown in Figure 1, one can place 5 paths in an $\mathcal{H}$-graph whose conflict graph is a 5-cycle. Thus we have

Lemma 3.6 *Let G be a directed graph, S a sufficient set for G, and K an $\mathcal{H}$-graph in G with characteristic path P. Then $S \cap P \neq \phi$.*

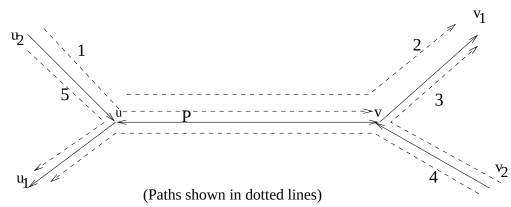

Figure 1: An $\mathcal{H}$-graph

First let us show that $\mathcal{H}$-graphs are the only structures that necessitate converters when $s(G)$ is a tree.

Lemma 3.7 *Suppose $s(G)$ is a tree. Then the empty set is sufficient for G if and only if G has no $\mathcal{H}$-graph.*

Proof. We know that if the empty set is sufficient, then G has no $\mathcal{H}$-graph. We prove the converse by induction on the number of edges of G. An edge (u, v) will be called *internal* if neither u nor v is a leaf.

First suppose G has an internal edge (u, v) for which $(v, u) \notin E$. Let $\mathcal{P}$ be a set of paths. If we delete the edge (u, v), $s(G)$ is partitioned into two components, R_1 and R_2. Construct two subgraphs of G: G_1, equal to the union of R_1 and the edge (u, v); and G_2, equal to the union of R_2 and (u, v). Since (u, v) is an internal edge, both these subgraphs have fewer edges than G. Restrict $\mathcal{P}$ to each subgraph as follows: if there is a path through (u, v), we only retain the edges of the path that lie in the respective subgraph. Paths not using (u, v) remain unchanged. Now, by induction hypothesis, we can color the paths in G_1 and G_2 with at most $\nu(\mathcal{P})$ colors. We can now merge these two subgraphs by permuting the labels of colors used in G_1 so that paths using (u, v) receive the same color in both subgraphs.

Thus we may consider the case in which all internal edges come in bi-directional pairs. If G contains at most one branching node, then it is a subgraph of a spider. Hence we may further focus on the case in which G has at least two branching nodes; if we consider two branching nodes u, v at minimum distance in G, then the path P between them is bi-directional and consists only of straight nodes. Moreover, by our assumption that G has no

$\mathcal{H}$-graph, P has the property that at one end, say u, all other incident edges come in the same direction. We now conclude by induction as follows. As shown in the figure below, we create two copies of u: u' and u''. We replace P by two bi-directional paths - from u' to v and from v to u''. Moreover, all incoming edges at u are now incident on u' (see figure).

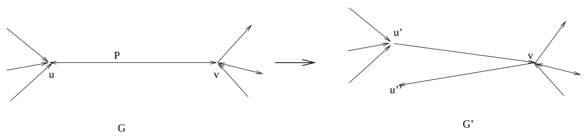

Figure 2: Replacing a bi-directional path by two directed paths

Let this new graph be G'. Note that G' can have paths (e.g., u' to v to u'') for which we may not have an analogue in G — but certainly, one can define a mapping from a path Q in G to a path Q' in G' in the usual manner — because no path Q in G will contain edges of P which go in both directions. Moreover, such a mapping clearly preserves the conflict graph and the load. So, if we can prove that empty set is sufficient for G', then empty set will be sufficient for G also. But G' has directed internal edges and so, by induction we have the result. ■

Now we consider the case in which $s(G)$ may contain cycles. A sufficient set need not meet every cycle in $s(G)$. However, our goal is to apply a sequence of transformations to G so that it has the following property: a set of nodes is sufficient if and only if it meets each undirected cycle of $s(G)$ and the characteristic path of each $\mathcal{H}$-graph. To do this, we define the following notions.

We say that a node of G is a *converging point* if it has at least two incident edges, so that each is oriented into v, or each is oriented out of v. We define a *bounded path* in G as follows. Let P be a bi-directional path with ends u and v. Let the other edges incident at u be called E_u and those incident at v be called E_v. Suppose all edges in E_u and E_v are directed into nodes u and v (respectively), or all are directed out of nodes u and v. Further, suppose there are no edges except those of P incident on any internal vertex of P. Then P is called a *bounded path* of G (see figure below). We say that a graph G is *robust* if it contains

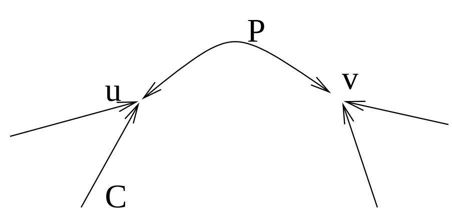

Figure 3: Example of a bounded path (P in the figure)

no converging points and no bounded paths.

The overview of our approach is now as follows.

(1) We apply a sequence of transformations to G, obtaining a robust graph G' with an "equivalent" minimum sufficient set problem.

(2) We show that in a robust graph, a set of nodes is sufficient if and only if it meets each undirected cycle of $s(G)$ and the characteristic path of each $\mathcal{H}$-graph.

(3) We formulate the problem of meeting all undirected cycles and characteristic paths as a *generalized feedback set* problem, which we approximate via the standard undirected feedback set problem.

Transforming G into a robust graph. Suppose v is a converging point of G, with incident edges $(v, u_1), (v, u_2), \ldots, (v, u_d)$. (The case in which all edges are incoming is strictly analogous.) Define a new graph G_v as follows: G_v is same as G except that we split v into d parts, each joined to one of u_i's. More formally, $G_v = (V_v, E_v)$, where

$$V_v = (V - \{v\}) \cup \{v_{u_1}, v_{u_2}, \ldots, v_{u_d}\},$$

$$E_v = (E_v - \{(v, u_1), \ldots, (v, u_d)\}) \cup \{(v_{u_1}, u_1), \ldots, (v_{u_d}, u_d)\}.$$

We call this operation *splitting* of G at v.

If v is a converging point, then any path passing through v must have one of its end points on v; so we do not need to do wavelength conversion at v. Thus we have

Lemma 3.8 *No minimal sufficient set of G contains a converging point. Hence, if v is a converging point, a set S of vertices is a minimal sufficient set for G if and only if it is a minimal sufficient set for G_v.*

Our other transformation removes bounded paths from G, based on the following lemma.

Lemma 3.9 *Let P be a bounded path in G, and S a minimal sufficient set. Then S does not contain any internal vertex of P.*

Thus we can apply the following transformation to a graph with a bounded path P, preserving the structure of the minimum sufficient set. Suppose P has ends u and v, and the edges at u and v are incoming edges - the other case is exactly similar. Then, we can break the bounded path as follows: Split u and v (and every other vertex on P) into two vertices each - u', u'' and v', v''. Replace P by two paths: P' - a directed path from u' to v' and P'' - a directed path from v'' to u'' - let this new graph be called G'. Then it is clear that we can have 1-1 correspondence (in the usual way) between the paths of G and G'. In fact conflict graphs of two corresponding set of paths will also remain the same (see the figure below).

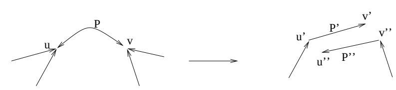

Figure 4: Breaking a bounded path

By applying a sequence of the above two types of transformations to G initially, we may assume for the remainder of this section that G is robust.

Sufficient sets in a robust graph. If K and K' are $\mathcal{H}$-graphs in G, we say that K' *encloses* K if the characteristic path of K is a subset of the characteristic path of K'. Enclosure, defined in this way, imposes a partial order on the $\mathcal{H}$-graphs of G, and we say an $\mathcal{H}$-graph is *minimal* if it is minimal with respect to enclosure.

Lemma 3.10 *(i) If K is a minimal $\mathcal{H}$-graph in G, with characteristic path P, then there is no edge of $G \setminus K$ incident on an internal node of P.*

(ii) The characteristic paths of two minimal $\mathcal{H}$-graphs are edge-disjoint.

A *minimal cycle* in an undirected graph is a cycle such that no two vertices on it are joined by an edge other than those on the cycle itself. Clearly, a graph has a cycle if and only if it has a minimal cycle. We can now prove the following lemma; the proof is fairly lengthy, and we omit it from this version. Essentially it constructs, for any minimal undirected cycle C in $s(G)$, a set of paths on the directed edges of C whose conflict graph is an odd cycle.

Lemma 3.11 *Let $G = (V, E)$ be robust, and C a minimal cycle of $s(G)$. Then any sufficient set must contain a vertex of C.*

A generalized feedback set problem. We are now ready to prove the following characterization theorem.

Theorem 3.12 *Let G be robust. Then a set of nodes S in G is sufficient if and only if it intersects each minimal cycle of $s(G)$, and the characteristic path of each minimal $\mathcal{H}$-graph.*

Proof. Lemmas 3.6 and 3.11 imply that S must meet each minimal cycle and the characteristic path of each minimal $\mathcal{H}$-graph. Conversely, suppose that S has this property, and consider the graph G' obtained by *splitting* G (as above) at each node in S. Each component X of G' has the property that it contains no $\mathcal{H}$-graph, and $s(X)$ is a tree. The result now follows from Lemma 3.7. ■

By Lemma 3.10, the characteristic paths of the minimal $\mathcal{H}$-graphs in G form a collection of edge-disjoint bi-directional paths $\{P_i\}$, each of whose internal vertices have degree 2 in $s(G)$. Given any sufficient set S for G, we can therefore transform it into a sufficient set S' of no greater size that meets each P_i at one of its ends.

Let G' denote the graph obtained from G by replacing each of the P_i by a single bi-directed edge e'_i between its endpoints, and deleting the internal nodes. As a consequence of Theorem 3.12 and the observation of the previous paragraph, our problem is equivalent to finding a set of minimum size that intersects each cycle in $s(G')$ and each of the edges e'_i.

We can phrase this as the following optimization problem.

> Generalized Feedback Set: Given an undirected graph $G = (V, E)$ and a set of edges $I \subseteq E$, find a set $S \subseteq V$ of minimum size such that S intersects all cycles of G and all edges in I.

However, this problem can be directly reduced to the standard feedback set problem: we simply replace each edge $e = (u, v)$ in I by a pair of parallel edges with endpoints u and v. This forces one of u or v to be in the output set S. (To avoid parallel edges, we can place a new vertex on each of the new edges). Since there is a polynomial-time 2-approximation for the undirected feedback set problem, we have

Corollary 3.13 *There is a polynomial-time* 2*-approximation for the minimum sufficient set problem in general directed graphs.*

4 Sufficient Sets and Feedback Sets in Planar Graphs

The transformations of the previous section preserve planarity, and hence to obtain a polynomial-time approximation scheme for the minimum sufficient set problem in planar graphs, we need only obtain one for the feedback set problem.

We develop our approximation scheme according to the following outline. First we show how to handle planar graphs of bounded degree, via an approach based on the Lipton-Tarjan planar separator theorem. We then identify a particular type of structure, a *highly connected pair*, and show that if this structure is not present, then a simple extension of the algorithm for bounded-degree graphs can be used. Finally we show how, given a highly connected pair, we can decompose the graph and solve the problem in a recursive fashion. We will continue to assume that G is connected.

The bounded-degree case. Let F^* denote an optimum feedback set, of size f^*. First, no cut-edge can take part in a cycle, so we may delete all cut-edges and henceforth assume that G is 2-edge-connected. Second, F^* need not contain any nodes of degree 2; thus, for each degree-2 node u with neighbors v and w, we may remove u and join v and w by an edge. By applying this operation repeatedly, we may assume that G has minimum degree 3. This operation may result in a graph G with parallel edges, even if the original graph did not have any; however, this will not pose a problem in the analysis below. (Note a feedback set is required to meet all two-edge cycles.)

By the Lipton-Tarjan planar separator theorem [18, 19], it is possible to remove a set S of $O(n/\sqrt{\log n})$ nodes from G, so that $G \setminus S$ has no component of size greater than $\log n$. By finding an optimal feedback set in each component, and adding the separator S to it, we obtain a feedback set of size at most $f^* + O(n/\sqrt{\log n})$. Unfortunately, f^* may be much smaller than $n/\sqrt{\log n}$, which renders the resulting approximation guarantee enormous.

However, as a first important step, we can show that when G has bounded maximum degree (and no cut-edges, and minimum degree 3), f^* must be large. Thus, in this case, the above approach provides an approximation scheme.

Theorem 4.1 *Let G be an n-node planar graph with minimum degree* 3*, maximum degree Δ, and no cut-edge. Then $f^* = \Omega(n/\Delta^2)$.*

Proof. Let $Z_1, Z_2, \ldots, Z_t$ denote the vertex sets of faces of G. By Euler's formula, $t = 2 + m - n$, where m denotes the number of edges of G; since G has minimum degree 3, we have $m \geq \frac{3}{2}n$, and hence $t \geq \frac{1}{2}n$.

Now, for any Z_i, the subgraph $G[Z_i]$ is not necessarily a simple cycle. However, it follows from the 2-edge-connectivity of G that $G[Z_i]$ contains a simple cycle. Hence, to prove the theorem, we need only exhibit a collection of $\Omega(n/\Delta^2)$ pairwise disjoint sets among $Z_1, Z_2, \ldots, Z_t$.

Consider the graph H whose vertex set is $\{Z_1, Z_2, \ldots, Z_t\}$ and with an edge between Z_i and Z_j if and only if $Z_i \cap Z_j \neq \phi$. Since each node v of G is incident to at most Δ faces in G, it contributes at most Δ^2 edges to H; hence, H has at most $\Delta^2 n \leq 2\Delta^2 t$ edges. This implies that at least $\frac{1}{2}t$ of the nodes in H have degree at most $8\Delta^2$. Let H' denote the subgraph of H induced on the nodes of degree at most $8\Delta^2$; this graph can be colored with at most $8\Delta^2$ colors, and hence it contains an independent set of size at least $t/(16\Delta^2) \geq n/(32\Delta^2)$. This independent set thus corresponds to a set of at least $n/(32\Delta^2)$ pairwise disjoint faces in G, as desired. ∎

Highly connected pairs. In view of Theorem 4.1, a natural algorithm for general planar graphs would be the following.

```
Algorithm 1:
  Define k = log^{1/5} n.
  Let X be the set of all vertices of degree at least k.
  Use the Lipton-Tarjan approach to find a feedback set Y in G\X.
  Output X ∪ Y.
```

Unfortunately, this algorithm does not necessarily provide a $(1+\varepsilon)$-approximation for sufficiently small $\varepsilon > 0$. We now define a type of subgraph that causes Algorithm 1 to perform badly, and prove that it is an approximation scheme when this type of subgraph is not present.

Let us fix a plane embedding of G. Suppose we have also fixed $\varepsilon > 0$, and seek a $(1+\varepsilon)$-approximation. Let r be a constant chosen to be sufficiently large relative to ε^{-1}: $r > 8\varepsilon^{-1}+3$ will suffice. We say that a vertex has *high degree* if its degree is at least $k = \log^{1/5} n$. We say that a high-degree vertex v is *highly connected* to a vertex u (or that (u, v) is a highly connected pair) if the following holds.

(1) There is a set of at least r node disjoint paths from v to u, denoted $P_1, \ldots, P_{r'}$ $(r' \geq r)$.

(2) For each pair $i \neq j$, any path connecting an internal node of P_i to an internal node of P_j must pass through either u or v.

We now show that if G does not contain any highly connected pairs, then Algorithm 1 is a $(1+\varepsilon)$-approximation. First, for a high-degree vertex v of degree d, consider the subgraph of G consisting of all vertices and edges incident to faces on which v is incident. In a standard way, we can decompose this subgraph into at least $d/2$ cycles; see the figure below. (Note that the cycles are not edge-disjoint.) Let $\mathcal{C}_v$ denote this set of cycles. (This can be defined formally as a set of those cycles containing v which enclose disjoint connected regions when edges of G are removed from the plane embedding.)

We now have the following fact.

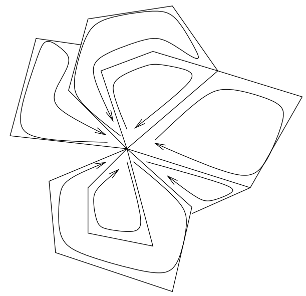

Figure 5: The set of cycles $\mathcal{C}_v$

Lemma 4.2 *If any vertex $u \neq v$ intersects more than r cycles in $\mathcal{C}_v$, then v is highly connected to u.*

Using this lemma, we can prove the following.

Lemma 4.3 *Suppose G has no highly connected pairs. Let X be a set of s high-degree vertices $v_1, \ldots, v_s$ with the property that $F^* \cap X = \phi$. Then $|F^*| \geq cs \log^{1/5} n$, where c is a constant depending on r.*

Proof. For a cycle C in $\cup_{i=1}^s \mathcal{C}_{v_i}$, define its *representative vertex* as that vertex from X which lies on C and has the smallest index among all the vertices of X lying on C. First we prove that the cardinality of $\cup_{i=1}^s \mathcal{C}_{v_i}$ is large.

Claim 4.4 $|\cup_{i=1}^s \mathcal{C}_{v_i}| \geq sk/2 - 3sr$ *(recall that r stands for a large constant).*

Proof. $\mathcal{C}_{v_i}$ and $\mathcal{C}_{v_j}$ contain a common cycle iff both v_i and v_j lie on this cycle. So, from Lemma 4.2, $\mathcal{C}_{v_i}$ and $\mathcal{C}_{v_j}$ can share at most r common cycles.

Construct a multigraph $G' = (V', E')$ as follows: $V' = \{v_1, \ldots, v_s\}$. For each cycle C in $\cup_{i=1}^s \mathcal{C}_{v_i}$, draw the following edges in G': suppose the representative vertex of C is v_{j_1} and the set $C \cap V' = \{v_{j_1}, v_{j_2}, \ldots, v_{j_p}\}$, where $j_1 < j_2 < \ldots < j_p$. Then join v_{j_1} to each of $v_{j_2}, \ldots, v_{j_p}$. Note that this may result in multiple edges between two vertices, but by Lemma 4.2, there can be at most r edges between any 2 vertices.

We claim that G' is planar. This is so because given the embedding of G, the edges drawn for each cycle C in $\cup_{i=1}^s \mathcal{C}_{v_i}$ remain inside C (since they all originate from the representative vertex, they do not cross) and the region enclosed by different cycles are disjoint, so no two edges cross.

Since G' is planar, if we replace all multiple edges by single edges, then it will have at most $3s$ edges. Since each single edge can account for at most r multiple edges, G' has at most $3sr$ edges.

Let n_{ij} denote the number of edges between v_i and v_j in G' (so $n_{ij} \leq r$). The total number of cycles in $\cup_{i=1}^s \mathcal{C}_{v_i}$ can be computed in the following manner

$$|\cup_{i=1}^s \mathcal{C}_{v_i}| \;=\; \sum_{i=1}^{s} |\mathcal{C}_{v_i} - (\cup_{j=1}^{i-1} \mathcal{C}_{v_j})|$$

Consider a cycle C contained in both the sets $\mathcal{C}_{v_i}$ and $(\cup_{j=1}^{i-1}\mathcal{C}_{v_j})$. There will be an edge between v_i and one of $v_1, \ldots, v_{i-1}$ in G' for this cycle C, because the representative vertex of C will have index smaller than i. Thus, $\mathcal{C}_{v_i} \cap (\cup_{j=1}^{i-1}\mathcal{C}_{v_j})$ has cardinality $\sum_{j=1}^{i-1} n_{ji}$. Thus,

$$|\cup_{i=1}^{s} \mathcal{C}_{v_i}| = \sum_{i=1}^{s}(|\mathcal{C}_{v_i}| - \sum_{j=1}^{i-1} n_{ji})$$

which gives the desired result. ■

F^* must intersect each of the cycles in $\cup_{i=1}^{s}\mathcal{C}_{v_i}$. Any vertex in F^* can intersect at most r cycles in $\mathcal{C}_{v_i}$ for any particular i. We construct a graph $G' = (V', E')$ where V' is the union of $\{v_1, \ldots, v_s\}$ and the vertices in F^*. For each cycle $C \in \cup_{i=1}^{s}\mathcal{C}_{v_i}$, let v_i be the representative vertex of C and u be a vertex in F^* which lies on C. We join v_i and u by an edge. Thus we have one edge in G' for each such cycle.

We claim that G' is planar. Each edge of G' can be constructed by drawing a curve between two vertices lying on a unique cycle in $\cup_{i=1}^{s}\mathcal{C}_{v_i}$. Now, the connected components of the plane defined by the cycles in $\cup_{i=1}^{s}\mathcal{C}_{v_i}$ are disjoint. So, none of the curves that we construct will intersect. Since G' is planar and each edge appears with multiplicity at most r, G' has at most $3(|F^*| + s)r$ edges. Combining this fact with Claim 4.4, we have $3(|F^*| + s)r \geq sk/2 - 3sr$, which implies the lemma. ■

Lemma 4.3 implies that we can add all high-degree vertices to F^* without affecting the approximation ratio by more than a $o(1)$ factor. Thus we have

Theorem 4.5 *If G has no highly connected pairs, then Algorithm 1 is an approximation scheme.*

Decomposing along a highly connected pair. We now consider the general case, when G may contain highly connected pairs. We first develop a type of *nesting structure* on these pairs that will be useful below.

Let (u, v) be a highly connected pair, with vertex-disjoint paths $P_1, \ldots, P_{r'}$ as in the definition, indexed from left to right in the plane embedding of G. Consider the subgraph R_{uv} consisting of u, v, and all nodes reachable from $P_2 \cup \cdots \cup P_{r'-1}$ by paths not passing through u or v. We refer to R_{uv} as the *region enclosed* by the highly connected pair (u, v), with u and v its *endpoints*; see the figure below.

The regions enclosed by highly connected pairs satisfy the following *nesting property*: for any two such regions, either one contains the other, or they are disjoint except possibly for one of their endpoints. We say that a highly connected pair is *minimal* if the region it encloses is minimal with respect to this containment relation. Based on this structure, we can now give a $(1+\varepsilon)$-approximation algorithm for finding a feedback set in a general planar graph. The algorithm is specified in Figure 6. It makes use of the following "fixed-parameter tractability" result (see Downey and Fellows [10]): for each fixed k, one can decide in time $O(n)$ whether a graph has a feedback set of size at most k.

Let F denote the set output by this algorithm. The fact that F is a feedback set for G is a consequence of the following lemma, which follows directly from the definition of highly connected pairs.

```
Algorithm 2:
   1.Define C = 4ε^{-1}.
   2.If G has no highly connected pairs, invoke Algorithm 1 and terminate.
   3.Let (u,v) be a minimal highly connected pair in G, enclosing a region R.
          Using the fixed-parameter feedback set algorithm [10],
          decide if the optimum feedback set K of R has size ≤ 2C.
   4.If |K| > 2C then
     (a)Let S denote R \ {u,v}.
          Use Algorithm 1 to find a (1+ε/2)-approximate feedback set F_0 for G[S].
              (Note that G[S] has no highly connected pair.)
     (b)Recursively compute a feedback set F' for G \ R.
     (c)Output F_0 ∪ F' ∪ {u,v} and terminate.
   5.Else |K| ≤ 2C.
     Compute the following three quantities, each in O(n) time [10].
       n_1:  the cardinality of the minimum feedback set F_1 of R containing both u and v.
       n_2:  the cardinality of the minimum feedback set F_2 of R containing u but not v.
       n_3:  the cardinality of the minimum feedback set F_3 of R containing v but not u.
   6.Three cases arise:
     (a)n_1 ≤ n_2, n_3:
        Recursively compute a feedback set F' for G \ R.
        Output F_1 ∪ F'.
     (b)n_2 < n_1, n_3 (the case when n_3 < n_1, n_2 is strictly analogous):
        Let S = R \ {v}.
        Recursively compute a feedback set F' for G \ S.
        Output F_2 ∪ F'.
     (c)n_2 = n_3 < n_1:
        Let S' = R \ {u,v}.
        Let G' be the graph obtained by removing S' from G and
           adding two new edges between u and v.
        Recursively compute a feedback set F' for G'.
           (Note that F' contains at least one of u or v.)
        If u ∈ F' (the case v ∈ F' is analogous), output F_2 ∪ F'.
```

Figure 6: The approximation algorithm for planar feedback set

Lemma 4.6 *Let (u,v) be a highly connected pair enclosing a region R. Let $w \in R$ and $w' \notin R$, $w, w' \neq u, v$. Then, any cycle containing w and w' must contain both u and v.*

Theorem 4.7 *For a given $\varepsilon > 0$, Algorithm 2 finds a feedback set in a planar graph of size at most $1 + \frac{\varepsilon}{2} + \frac{2}{C}$ times optimal. For $C = \frac{4}{\varepsilon}$, this is a $(1+\varepsilon)$-approximation.*

Proof. Recall that F^* denotes an optimum feedback set for G; for a set $X \subseteq V$, let F^*_X denote $F^* \cap X$. The proof is by induction on the number of highly connected pairs in G. When G has no highly connected pairs, the result is a consequence of Theorem 4.5. Otherwise, let (u,v) be the minimal highly connected pair considered in Step 3 of the algorithm, and K an optimum feedback set of the region R that it encloses.

First consider the case in Step 4, when $|K| > 2C$. Now by induction we have $|F'| \leq (1+\frac{\varepsilon}{2}+\frac{2}{C})|F^*_{G\setminus R}|$; since $|K| > 2C$, we also have $|F_0 \cup \{u,v\}| \leq (1+\frac{\varepsilon}{2})|F^*_R|+2 \leq (1+\frac{\varepsilon}{2}+\frac{2}{C})|F^*_R|$. Hence

$$|F| \quad = \quad |F_0 \cup \{u,v\}| + |F'|$$

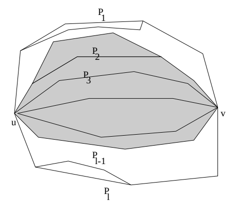

Figure 7: The region enclosed by (u, v) shown as the shaded region

$$\begin{aligned} &\le (1+\frac{\varepsilon}{2}+\frac{2}{C})(|F^*_{G\setminus R}|+|F^*_R|) \\ &= (1+\frac{\varepsilon}{2}+\frac{2}{C})|F^*|. \end{aligned}$$

Next, consider the case in Steps 5 and 6, when $|K| \le 2C$. We claim first that F^* must contain either u or v. For suppose not. Then F^* must intersect at least $r-3$ vertex disjoint paths between u and v, because there are at least $r-2$ such paths between u and v in R and any two of these form a cycle. Since $r > 2C+3$, this is a contradiction.

We now discuss the cases in Step 6 individually.

- Case 6(a): $|F_1| \le |F^*_R|$ because F^*_R must contain either u or v and $n_2, n_3 \ge n_1$. So

$$\begin{aligned} |F| &= |F_1| + |F'| \le |F^*_R| + (1+\frac{\varepsilon}{2}+\frac{2}{C})|F^*_{G\setminus R}| \\ &\le (1+\frac{\varepsilon}{2}+\frac{2}{C})|F^*|. \end{aligned}$$

- Case 6(b): Two cases arise :

(i) F^*_R contains u, but it doesn't contain v. In Step 6(b), S is defined as $R \setminus \{v\}$. $|F^*| = |F^*_R| + |F^*_{G\setminus S}|$ since the only vertex common between R and $G \setminus S$ is v. So

$$\begin{aligned} |F| &= |F_2| + |F'| \le |F^*_R| + (1+\frac{\varepsilon}{2}+\frac{2}{C})|F^*_{G\setminus S}| \\ &\le (1+\frac{\varepsilon}{2}+\frac{2}{C})|F^*|. \end{aligned}$$

(ii) F^*_R contains v; in this case v will be counted twice in F^*_R and $F^*_{G\setminus S}$. But observe that $|F^*_R| \ge |F_2| + 1$, since n_2 is strictly less than n_1 and n_3. Thus

$$\begin{aligned} |F| &= |F_2| + |F'| \\ &\le |F^*_R| - 1 + (1+\frac{\varepsilon}{2}+\frac{2}{C})|F^*_{G\setminus S}| \\ &\le (1+\frac{\varepsilon}{2}+\frac{2}{C})(|F^*_R| + |F^*_{G\setminus S}| - 1) \\ &= (1+\frac{\varepsilon}{2}+\frac{2}{C})|F^*|. \end{aligned}$$

- Case 6(c): Suppose without loss of generality that F' contains u. Note that $|F| = |F_2| + |F'| - 1$ since u is counted twice. Two cases arise :

(i) F_R^* contains exactly one of u and v. In this case, we have one point in common between F_R^* and $F_{G\setminus S'}^*$ where S' was defined as $R \setminus \{u, v\}$. (Note that the new edges cause no problem since we've already proved that F^* contains either u or v.) Thus,

$$\begin{aligned} |F| &= |F_2| + |F'| - 1 \\ &\leq |F_R^*| + (1 + \frac{\varepsilon}{2} + \frac{2}{C})|F_{G\setminus S'}^*| - 1 \\ &\leq (1 + \frac{\varepsilon}{2} + \frac{2}{C})(F_R^* + |F_{G\setminus S'}^*| - 1) \\ &= (1 + \frac{\varepsilon}{2} + \frac{2}{C})|F^*|. \end{aligned}$$

(ii) F_R^* contains both u and v. In this case, $|F_2| \leq |F_R^*| - 1$ (since $n_2 < n_1$). Also, $|F^*| = |F_R^*| + |F_{G\setminus S}^*| - 2$. Thus

$$\begin{aligned} |F| &= |F_2| + |F'| - 1 \\ &\leq |F_R^*| - 1 + (1 + \frac{\varepsilon}{2} + \frac{2}{C})|F_{G\setminus S}^*| - 1 \\ &\leq (1 + \frac{\varepsilon}{2} + \frac{2}{C})(|F_R^*| + |F_{G\setminus S}^*| - 2) \\ &= (1 + \frac{\varepsilon}{2} + \frac{2}{C})|F^*|. \qquad \blacksquare \end{aligned}$$

Acknowledgements. We thank Nathan Edwards and S. Keshav for valuable discussions.

References

[1] A. Aggarwal, A. Bar-Noy, D. Coppersmith, R. Ramaswami, B. Schieber, M. Sudan, "Efficient routing and scheduling algorithms for optical networks," *Proc. 5th ACM-SIAM Symp. on Discrete Algorithms*, 1994.

[2] V. Auletta, I. Caragiannis, C. Kaklamanis, P. Persiano, "Bandwidth allocation algorithms on tree-shaped all-optical networks with wavelength converters," *Proc. SIROCCO*, 1997.

[3] Y. Aumann, Y. Rabani, "Improved bounds for all-optical routing," *Proc. 6th ACM-SIAM Symp. on Discrete Algorithms*, 1995, pp. 567–576.

[4] V. Bafna, P. Berman, T. Fujito, "Constant factor approximation of the weighted feedback vertex set problem," *Proc. ISAAC*, 1995.

[5] B. Baker, "Approximation algorithms for NP-complete problems on planar graphs," *Journal of the ACM*, 41(1994), pp. 153–180.

[6] R. Bar-Yehuda, D. Geiger, J. Naor, R. Roth, "Approximation algorithms for the vertex feedback set problem with applications to constraint satisfaction and Bayesian inference," *Proc. 5th ACM-SIAM Symp. on Discrete Algorithms*, 1994.

[7] R. A. Barry, V. W. S. Chan, K. L. Hall, E. S. Kintzer, J. D. Moores, K. A. Rauschenbach, E. A. Swanson, L. E. Adams, C. R. Doerr, S. G. Finn, H. A. Haus, E. P. Ippen, W. S. Wong, M. Haner, "All-Optical Network Consortium - Ultrafast TDM Networks," *IEEE Journal on Selected Areas in Communications*, 14:5(1996).

[8] A. Becker, D. Geiger, "Approximation algorithms for the loop cutset problem," *Proc. 10th Conference on Uncertainty in AI*, 1994.

[9] F. Chudak, M. Goemans, D. Hochbaum, D. Williamson, "A primal-dual interpretation of two 2-approximation algorithms for the feedback vertex set problem in undirected graphs," *Operations Research Letters*, to appear.

[10] R. Downey, M. Fellows, "Parametrized Computational Feasibility," in *Feasible Mathematics II*, P. Clote and J. Remmel, eds., Birkhauser, 1994.

[11] M. Goemans, D. Williamson, "Primal-dual approximation algorithms for feedback problems in planar graphs," *Proc. Integer Programming and Combinatorial Optimization*, 1996.

[12] P.E. Green, *Fiber-Optic Communication Networks*, Prentice-Hall, 1993.

[13] S. Keshav, *An Engineering Approach to Computer Networking*, Addison-Wesley, 1997.

[14] J. Kleinberg, É. Tardos, "Approximations for the Disjoint Paths Problem in High–Diameter Planar Networks," *Proc. 27th ACM Symp. on Theory of Computing*, 1995, pp. 26–35.

[15] M. Kovacevic, A. Acampora, "On the benefits of wavelength translation in all-optical clear-channel networks," *Proc. IEEE INFOCOM*, 1995.

[16] V. Kumar, E. Schwabe, "Improved access to optical bandwidth," *Proc. 8th ACM-SIAM Symp. on Discrete Algorithms*, 1997.

[17] K.C. Lee, V. Li, "Routing and switching in a wavelength convertible optical network," *Proc. IEEE INFOCOM*, 1993.

[18] R. Lipton, R. Tarjan, "A separator theorem for planar graphs," *SIAM J. Applied Math.* 36(1979).

[19] R. Lipton, R. Tarjan, "Applications of a planar separator theorem," *SIAM J. Computing*, 9(1980).

[20] M. Mihail, C. Kaklamanis, S. Rao, "Efficient access to optical bandwidth," *Proc. 36th IEEE Symp. on Foundations of Computer Science*, 1995.

[21] R. Pankaj, *Architectures for Linear Lightwave Networks*, Ph.D. Thesis, MIT Dept. of EECS, 1992.

[22] C. Partridge, *Gigabit Networking*, Addison-Wesley, 1993.

[23] Y. Rabani, "Path-coloring on the mesh," *Proc. 37th IEEE Symp. on Foundation of Comp. Sci.*, 1996.

[24] P. Raghavan, E. Upfal, "Efficient all-optical routing," *Proc. 26th ACM Symp. on Theory of Computing*, 1994.

[25] R. Ramaswami, G. Sasaki, "Multiwavelength optical networks with limited wavelength conversion," *Proc. IEEE INFOCOM*, 1997.

[26] R. Tarjan, "Decomposition by clique separators," *Discrete Math.*, 55(1985), pp. 221–232.

[27] G. Wilfong, P. Winkler, "Ring routing and wavelength translation," *Proc. 9th ACM-SIAM Symp. on Discrete Algorithms*, 1998.