On shortest-path all-optical networks without wavelength conversion requirements

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Thomas Erlebach, Stamatis Stefanakos

On Shortest-Path All-Optical Networks without Wavelength Conversion Requirements

TIK-Report
Nr. 153, October 2002

Institut für Technische Informatik und Kommunikationsnetze
Computer Engineering and Networks Laboratory

Thomas Erlebach, Stamatis Stefanakos
On Shortest-Path All-Optical Networks without Wavelength Conversion Requirements
October 2002
Version 2 (Revised: 14.11.2002)
TIK-Report Nr. 153

Computer Engineering and Networks Laboratory,
Swiss Federal Institute of Technology (ETH) Zurich

Institut für Technische Informatik und Kommunikationsnetze,
Eidgenössische Technische Hochschule Zürich

Gloriastrasse 35, ETH Zentrum, CH-8092 Zürich, Switzerland

On Shortest-Path All-Optical Networks without Wavelength Conversion Requirements*

Thomas Erlebach   Stamatis Stefanakos
Computer Engineering and Networks Laboratory (TIK)
ETH Zürich, CH-8092 Zürich, Switzerland
{erlebach|stefanak}@tik.ee.ethz.ch

Abstract

In all-optical networks with wavelength division multiplexing, every connection is routed along a certain path and assigned a wavelength such that no two connections use the same wavelength on the same link. For a given set $\mathcal{P}$ of paths (a routing), let $\chi(\mathcal{P})$ denote the minimum number of wavelengths in a valid wavelength assignment and let $L(P)$ denote the maximum link load. We always have $L(\mathcal{P}) \leq \chi(\mathcal{P})$. Motivated by practical concerns, we consider routings containing only shortest paths. We give a complete characterization of undirected networks for which any set $\mathcal{P}$ of shortest paths admits a wavelength assignment with $L(\mathcal{P})$ wavelengths. These are exactly the networks that do not benefit from the use of (expensive) wavelength converters if shortest-path routing is used. We also give an efficient algorithm for computing a wavelength assignment with $L(\mathcal{P})$ wavelengths in these networks.

1 Introduction

In all-optical networks that employ wavelength division multiplexing, a connection is established by first choosing a path from the sender to the receiver and then assigning a wavelength for that connection to all the links of the path. The wavelength assignment has to be done so that no two connections that share a link are transmitted through the same wavelength. Since the number of available wavelengths is limited, one is interested in minimizing the number of utilized wavelengths for a given set of connections. One technique that helps cut down the number of necessary wavelengths for operating a network is that of wavelength conversion. A wavelength converter is placed in some node of the network and has the ability of altering the transmitting wavelength of any incoming signal. An interesting algorithmic problem that arises then is that of placing as few converters as possible in suitable positions of the network in order to optimize its

*Research partially supported by the Swiss National Science Foundation under Contract No. 21-63563.00 (Project AAPCN) and the EU Thematic Network APPOL II (IST-2001-32007), with funding provided by the Swiss Federal Office for Education and Science (BBW).

capacity. In this paper we focus on networks where connections are always established through shortest paths and we characterize the networks that do not profit from the use of converters.

The network can be naturally modeled by a graph $G = (V, E)$, a connection in the network can be seen as a path on G, and wavelengths can be regarded as colors. If the network does not employ wavelength conversion then a wavelength assignment for a set $\mathcal{P}$ of paths is an assignment of a color to each path in $\mathcal{P}$. A valid coloring is one in which no two paths that use the same edge get assigned the same color. For a given set $\mathcal{P}$ of paths (a routing) we denote by $\chi(\mathcal{P})$ the minimum number of colors needed for a valid coloring of $\mathcal{P}$. A trivial lower bound for $\chi(\mathcal{P})$ is the congestion or load $L(\mathcal{P}) = \max_{e \in E} L_e(\mathcal{P})$ of the network, where $L_e(\mathcal{P})$ is the number of paths in $\mathcal{P}$ that use e. If the network has wavelength converters a coloring is an assignment of a color to every edge of each path. In this case, a valid coloring has to satisfy the additional constraint that the color assignments to two consecutive edges of a path can only differ if there is a converter between the two edges.

Obviously, even with the use of converters, we will always need at least $L(\mathcal{P})$ colors for a valid coloring of $\mathcal{P}$. However, the placement of converters in a network can reduce the number of wavelengths needed for a fixed set of connections. It can be the case, for example, that while there is no single wavelength, from the already used ones, available along a path there are different available wavelengths along parts of that path. If there are converters in suitable positions along that path then these different wavelengths can be exploited in order to serve the connection which otherwise would have had to be assigned a new wavelength.

Ideally, by placing converters in some nodes of the network we can guarantee that any routing $\mathcal{P}$ can be accommodated with $L(\mathcal{P})$ wavelengths. This is easily seen to be the case if all nodes of the network are equipped with wavelength converters. However, the cost of these devices is prohibitive for such improvident use. Therefore, the network designer is typically interested in equipping only a small subset of the network terminals with converters, while still achieving the same capacity usage as if there where converters everywhere. Motivated by this, Wilfong and Winkler [11] introduced the following problem:

Minimum Sufficient Set. *Given a graph $G = (V, E)$, find a sufficient set S for G, i.e., a set $S \subseteq V$ such that any set $\mathcal{P}$ of paths on G can be colored with $L(\mathcal{P})$ colors if we place wavelength converters on the vertices of S. The goal is to minimize the size of S.*

Wilfong and Winkler [11] proved that Minimum Sufficient Set is $\mathcal{NP}$-hard even for planar bidirected graphs (a bidirected graph is a directed graph where $(u, v) \in E \Rightarrow (v, u) \in E$). Moreover they showed that the only bidirected graphs that admit the empty sufficient set are spiders, i.e., trees with at most one vertex of degree greater than two, and that rings admit a sufficient set of size 1. Finally, they described an efficient way of determining whether a set S is sufficient for a bidirected graph G: one modifies G by "exploding" each node $s \in S$ into degree-of-s-many copies, each of which is made adjacent to one of the old neighbors of s. S is sufficient for G if and only if every component of the graph obtained after this modification is a spider.

Extending this work, Kleinberg and Kumar [7] gave a 2-approximation algorithm for directed graphs and a polynomial time approximation scheme for directed planar graphs using techniques based on the undirected feedback vertex set problem. They also showed that any improvement on the approximation ratio for Minimum Sufficient Set on bidirected graphs would lead to a corresponding improvement for vertex cover. The approach of Kleinberg and Kumar can be extended to give a linear time algorithm for Minimum Sufficient Set in directed graphs of bounded treewidth [4].

1.1 Our Contribution

As described above, the previous work concentrated on the study of Minimum Sufficient Set in bidirected or directed graphs. These graphs serve as models for networks with unidirectional links, i.e., networks that support only one-way communication. In this paper, we turn to undirected graphs. Undirected graphs model networks where the physical links are bidirectional or networks with unidirectional links with the additional property that whenever a connection is established in one way, the reverse connection must also be established through the same path and must be transmitted over the same wavelength.

It is easy to see that the only undirected graphs that admit the empty sufficient set are chains. Furthermore, using the technique developed by Wilfong and Winkler, one can show that Minimum Sufficient Set is polynomial in undirected graphs: we simply have to place a converter in every node of degree greater than or equal to 3 or in any single node if the graph is a cycle (for more details see [4]). Nevertheless, such placement of converters is not satisfying. For example, in the case where the network is a clique we will need to place a converter in every node; however, it is unlikely that we would need any converter at all in practice since in a clique most connections would be carried over a single link. In order to capture this real-world scenario we restrict ourselves to shortest-path routings, i.e., we are interested in placing as few converters as possible so that any set $\mathcal{P}$ of shortest paths can be colored with $L(\mathcal{P})$ colors. More formally, we introduce the following problem:

Minimum SP-Sufficient Set. *Given a graph $G = (V, E)$, find an SP-sufficient set S for G, i.e., a set $S \subseteq V$ such that any set $\mathcal{P}$ of shortest paths on G can be colored with $L(\mathcal{P})$ colors if we place wavelength converters on the vertices of S. The goal is to minimize the size of S.*

We note that this problem is of significant practical importance since shortest-path routing is a common strategy in optical networks, see for example [12]. Moreover, it imposes a weaker condition than Minimum Sufficient Set (any sufficient set is an SP-sufficient set) and therefore can help decrease the cost of converters in the design of optical networks.

In this paper we give a complete characterization of the networks that admit the empty SP-sufficient set. We show that the block graph of such networks is a chain and all internal blocks are cliques. The outer blocks are a special case of co-bipartite graphs. If the graph consists of only one block we show that its diameter is less than or equal to 3. For the case of diameter 3 we prove that the graph belongs to a special class of co-bipartite graphs. For the case where the diameter is less than 3 we show that the graph admits the empty SP-sufficient set if and only if its edges can be 2-colored in a certain way. In all cases our proofs provide efficient algorithms for recognizing graphs that admit the empty SP-sufficient set and for optimally coloring shortest-path routings on these graphs. An interesting aspect of our work is that we do not need to consider complicated routings in order to obtain the characterization: all graphs that require a converter admit a witness routing $\mathcal{P}$ with $L(\mathcal{P}) = 2$ and $\chi(\mathcal{P}) = 3$ whose conflict graph is an odd cycle.

1.2 Other Related Work

Perhaps the main algorithmic problem that arises in optical networks is that of routing and wavelength assignment, i.e., given a set of connections in a network find a routing and a valid coloring of that routing so that the number of colors is minimized. Most research has focused on the case where the network does not have any conversion capabilities. In that case the problem is known to be $\mathcal{NP}$-hard even for simple topologies like rings and trees [2, 11], and in the case of rings, even if the routing is part of the input [5]. Consequently, a lot of effort has been put in designing approximation algorithms for specific network topologies, see e.g., [10, 9, 3].

There has been another line of research on wavelength converters of bounded degree. A converter of bounded degree does not have full conversion capabilities, i.e., it can transform color i to only a few other colors. We refer the reader to [1, 6] and the references therein.

1.3 Outline

The rest of the paper is structured as follows. In the following section we provide the necessary notation and give some insight to the techniques used in our proofs. In Section 3 we characterize graphs that do not need converters and contain cut-vertices while in Section 4 we turn to biconnected graphs. Finally, in Section 5 we summarize our results and discuss about future work.

2 Preliminaries

Throughout this paper, a graph $G = (V, E)$ is finite, simple, and undirected unless explicitly stated otherwise. For a graph G we will denote by $V(G), E(G)$ its vertex-set and edge-set respectively. In some cases, in order to simplify notation, we will write $v \in G$ instead of $v \in V(G)$. For $U \subseteq V$, $G[U]$ denotes the graph on U whose edges are the edges of G with both endpoints in U, i.e., $G[U]$ is the subgraph of G *induced* by U. The *distance* $d(u, v)$ in G of two vertices u, v is the number of edges in a shortest $u - v$ path in G (if no such path exists then $d(u, v) := \infty$). The *eccentricity* of a vertex v in G, $\mathrm{ecc}(v) = \max_{u \in V} d(u, v)$, is the maximum distance of v to all other vertices of G. The *diameter* of G, $\mathrm{diam}(G) = \max_{v \in V} \mathrm{ecc}(v)$, is the maximum eccentricity over all vertices of G. The degree, $\deg(v)$, of a vertex v in G is the number of edges incident to v. We denote by $\Delta(G) = \max_{v \in V} \deg(v)$ the maximum degree of G. We denote an induced cycle with k edges and k vertices by C^k and an induced path with k vertices and $k - 1$ edges by P^k. A maximal connected subgraph of G without a cut-vertex is a *block* of G. The *block graph* of G is the bipartite graph on $A \cup \mathcal{B}$ and edges aB for $a \in A$ and $B \in \mathcal{B}$ if $a \in B$, where A is the set of cut-vertices of G and $\mathcal{B}$ is the set of blocks of G. The *conflict graph* of a given set of paths is the graph with one vertex for each path and an edge between two vertices if the corresponding paths share an edge.

Obviously, a graph that does not need converters can not contain a configuration that allows us to construct a set $\mathcal{P}$ of shortest paths with $L(\mathcal{P}) < \chi(\mathcal{P})$. The configurations of this type that we will mainly use in this paper are shown in Fig. 1. In the text, in order to exhibit such a configuration we will refer to the vertices that induce the configuration; for example, we say that we have an antenna around $uvzwyx$ (Fig. 1(b)). For the case of the claw we say that we have a claw *at* x. All these configurations allow us to construct a set of shortest paths with load 2 whose conflict graph is an odd cycle and hence require 3 colors for a valid coloring. For example, in the case of the tent we can take v_1v_2u, v_2uv_5, uv_5v_6, $v_1v_2v_3$, $v_2v_3v_4$, $v_3v_4v_5$, and $v_4v_5v_6$. Notice that for some of

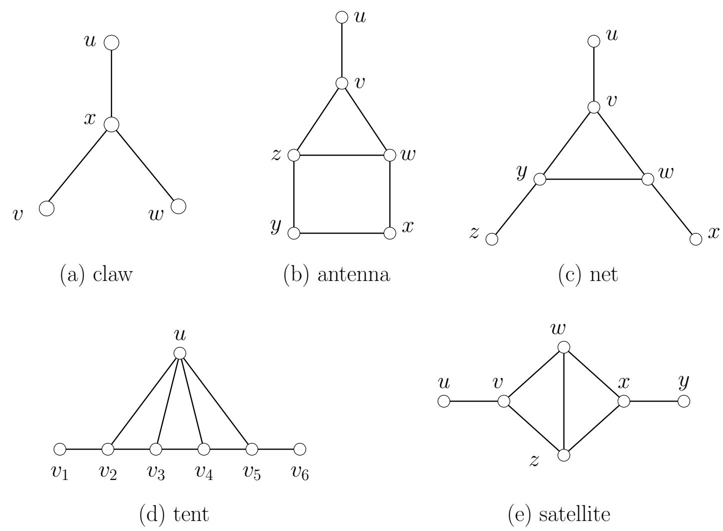

Figure 1: Some configurations that allow the construction of a set $\mathcal{P}$ of shortest paths with $L(\mathcal{P}) < \chi(\mathcal{P})$.

the configurations shown in Fig. 1, in order to find such a set of paths, we need to be able to find a shortest path of length 3 on them. This is the case for the net, the antenna and the satellite. For the latter for example, we can take the following shortest paths of length 2: uvw, vwx, wxy, yxz, xzv and zvu. In order to have an odd number of paths we need $d(u,x) = 3$ or $d(v,y) = 3$ so that we can exchange two paths of length 2 with one of length 3. Other configurations which allow us to construct a set $\mathcal{P}$ of shortest paths with $L(\mathcal{P}) < \chi(\mathcal{P})$ are induced cycles of length greater than or equal to 5. On an induced odd cycle we can just take all paths of length two. On an induced even cycle the set of all paths of length two has even cardinality and thus we need one path of length three.

In some of the proofs that follow, in order to exhibit the structure of the graphs that admit the empty SP-sufficient set, we will do case analyses regarding the presence or absence of some edges of the graph. In these case analyses we will refer to bad configurations, i.e., configurations that allow us to construct a set of shortest paths $\mathcal{P}$ with $L(\mathcal{P}) < \chi(\mathcal{P})$, using the terminology introduced above, even though some "future" edges that have not yet been considered might not allow the construction of such a set $\mathcal{P}$. For example, in Fig. 1(a), and assuming that we have eliminated the presence of edges uv, uw already by the case analysis, we say that we have a claw at x even though edge vw might be there. In this case this actually *forces* $vw \in E$ since otherwise this would give a bad configuration. Alternatively we also say that the claw can only be *broken* by vw.

3 Graphs with Cut-Vertices

First, assume that $G = (V, E)$ is connected but not biconnected and admits the empty SP-sufficient set.

Lemma 1 *The block graph of G is a chain.*

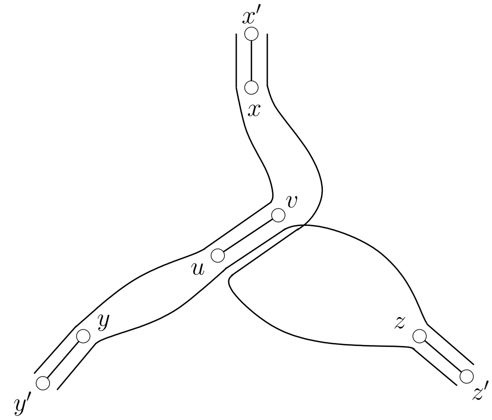

Figure 2: Proof of Lemma 1: not all three paths can use the same edge.

Proof. We will show that every block of G contains at most two cut-vertices and every cut-vertex is contained in exactly two blocks. For the first part assume for a contradiction that there exists a block C of G containing three cut-vertices x, y, z. Let x', y', z' be neighbors of x, y and z respectively outside C. Consider three shortest $x' - y', x' - z'$ and $y' - z'$ paths. The only edges outside C that are touched by these paths are xx', yy', zz' and all three paths pairwise intersect on these edges. Furthermore, not all three paths use the same edge. To see this assume to the contrary that there is an edge $e = uv$ used by all three paths (consider Fig. 2). Since uv is on a shortest $x' - y'$ path we have that $d(v, x') \neq d(u, x')$. Assume that $d(v, x') < d(u, x')$. Then, $d(u, y') < d(v, y')$ and also $d(u, z') < d(v, z')$ since uv is also on the $x' - z'$ path. But this is a contradiction since uv is on the $y' - z'$ shortest path. Hence, we can find three shortest paths with load 2 that all pairwise intersect and thus require three colors for a valid coloring, a contradiction. For the second part, notice that if there exists a cut-vertex contained in three or more blocks, we have a claw and hence we can find a set of shortest paths $\mathcal{P}$ with $L(\mathcal{P}) < \chi(\mathcal{P})$. □

Lemma 2 *Let C be a block of G containing a cut-vertex x. Then the following hold:*

(i) For all $u \in C$, $d(u, x) \leq 2$.

(ii) For all $u, v \in C$ such that u, v are not cut-vertices, $d(u, v) \leq 2$.

(iii) If C contains a second cut-vertex y then $d(x, y) = 1$.

Proof. (i) Let $u \in C$ and assume $d(u, x) > 2$. Let $xa_1 \ldots a_k u$, $xb_1 \ldots b_l u$ be two disjoint $x - u$ paths, each of length at least 3 such that $k + l$ is minimum, and let x' be a neighbor of x outside C. Notice that if there are edges $a_i b_{j'}, a_{i'} b_j$ with $i < i'$ and $j < j'$ then $i' = i + 1$ and $j' = j + 1$ since otherwise $k + l$ is not minimum. Also, if there are two such edges then we also have $a_i b_j, a_{i+1} b_{j+1} \in E$ since otherwise we would have a claw. We distinguish cases depending on the edges between a_1, b_2 and a_2, b_1. In all cases we have that $a_1 b_1 \in E$ because of a claw at x. If both $a_1 b_2, a_2 b_1 \in E$ then we have that $a_2 b_2 \in E$ because of claws at a_1 and b_1. Also, $a_1 b_3 \notin E$ because if this was not the case, $k + l$ would not be minimum. Hence, we have a satellite around $x' x a_1 b_1 b_2 b_3$ with a path of length 3 on it, say $x' x a_1 b_2$, a contradiction (the same holds if $b_3 = u$). If there are no edges between a_1, b_2

and a_2, b_1 and furthermore $a_2b_2 \notin E$ we have a net around $x'xa_1a_2b_1b_2$ with a path of length 3 on it, say $x'xa_1a_2$, a contradiction. If $a_2b_2 \in E$ we have an antenna around $x'xa_1a_2b_1b_2$ with a path of length 3 on it, say $x'xa_1a_2$, a contradiction. Consider the case where there is only one edge between a_1, b_2 and a_2, b_1, say a_1b_2. We have that $a_2b_2 \in E$ because of a claw at a_1. We have a satellite around $x'xa_1b_1b_2b_3$ with a path of length 3 on it, namely the path $x'xb_1b_2$. If $b_3 = u$ the satellite can not be broken and hence we reach a contradiction. Otherwise, we have that $a_1b_3 \in E$. We have a tent around $x'xa_1b_1b_2b_3b_4$. If $b_4 = u$ the tent can not be broken and we reach a contradiction; otherwise we have that $a_1b_4 \in E$ but this creates a claw at a_1 that can not be broken, a contradiction.

(ii) Let u, v, be two arbitrary vertices in C that are not cut-vertices and assume to the contrary that $d(u,v) > 2$. Let x' be a neighbor of x outside C. By (i) we have that $d(x,u) \leq 2$ and $d(x,v) \leq 2$. We distinguish cases according to the distance of u, v from x. We can not have $d(x,u) = d(x,v) = 1$ since then $d(u,v) \leq 2$, a contradiction to our assumption. Assume $d(x,v) = 2$ and let xav be an $x - v$ path of length 2. If $d(x,u) = 1$ we must have $au \in E$ because of a claw at x and hence $d(u,v) \leq 2$, a contradiction. For the case where $d(x,u) = 2$ we can assume that there exists an $x - u$ path of length 2, disjoint from xav, since otherwise we have that $d(u,v) \leq 2$. Let xbu be such a path. We have that $ab \in E$ because of a claw at x. This creates a net around $x'xbuav$ with a path of length 3 on it, say $x'xbu$, and hence we should have $au \in E$ or $bv \in E$ or $uv \in E$. In all cases we have that $d(u,v) \leq 2$, a contradiction.

(iii) Assume to the contrary that $d(x,y) \geq 2$. By (i) we have that $d(x,y) \leq 2$ and hence there exist two disjoint $x - y$ paths such that one of them is of length 2. Let xvy and $xa_1a_2 \ldots a_ky$ be two such paths with k minimum and let x', y' be neighbors of x, y respectively outside C. We reach a contradiction by exhibiting a set of paths $\mathcal{P}$ with $L(\mathcal{P}) = 2$ and $\chi(\mathcal{P}) = 3$. Consider the paths $x'xa_1$, $xa_1a_2, a_1a_2a_3, \ldots, a_{k-1}a_ky$, a_kyy', $x'xvy, vyy'$. All of them are shortest paths, their load is 2 and their intersection graph is a cycle. If k is odd, the cycle is of odd length and hence we need 3 colors for a valid coloring of $\mathcal{P}$. Otherwise replace the path $x'xvy$ with two paths of length 2. □

Lemma 3 *Let C be a block of G containing two cut-vertices x, y. For all $u \in C$, $d(u,x) = 1$ or $d(u,y) = 1$.*

Proof. Let u be an arbitrary vertex in C. Assume for a contradiction that neither x nor y are adjacent to u. From Lemma 2 we have that $d(u,x) \leq 2$, $d(u,y) \leq 2$ and $d(x,y) = 1$. Therefore, $d(u,x) = d(u,y) = 2$ and there exists a vertex w adjacent to both x and u. Since y is adjacent to x we have that $wy \in E$ since otherwise there would be a claw at x. Let x', y' be neighbors of x, y respectively outside C. We have a net around $x'xwuyy'$ with the path $x'xyy'$ of length 3 on it and no possibility to break it, a contradiction. □

Corollary 4 *Every block C of G containing two cut-vertices is a clique.*

Proof. Let C be a block of G containing two cut-vertices x, y. Consider an arbitrary vertex $u \in C$. From the previous lemma we have that u is adjacent to x or to y. Assume w.l.o.g. $ux \in E$. We must have $uy \in E$ since otherwise we get a claw at x. Hence, any vertex in C that is not a cut-vertex is adjacent to both x and y. Therefore, any two vertices in C are adjacent to each other since otherwise we would have claws at x and y. □

Lemma 5 *Let C be a block of G containing exactly one cut-vertex x and let N_1, N_2 be the subgraphs of G induced by the neighborhoods of x in C in distance 1, 2 respectively. The following hold:*

(i) Both N_1, N_2 are cliques.

(ii) There is no C^4 in C.

(iii) If there exists a vertex $w \in N_2$ adjacent to two vertices $u, v \in N_1$, then all vertices in N_2 are adjacent to u or to v.

Proof. For the first part, notice that if N_1 is not a clique we get a claw at x. Consider now two vertices $u, v \in N_2$ and assume $uv \notin E$. If u, v are adjacent to the same vertex of N_1 we have a claw at that vertex. If they are adjacent to different vertices we have a net (since N_1 is a clique) with a path of length 3 on it. We proceed to show the second part. Assume there is a C^4 in C. Since N_1, N_2 are cliques the cycle should contain two vertices, say a, b from N_1 and two vertices, say c, d from N_2. Let $abcd$ be the cycle and let x' be a neighbor of x outside C. We have an antenna around $x'xabcd$ with a path of length 3 on it, a contradiction. For the last part of the lemma assume that there exists a vertex $w \in N_2$ adjacent to two vertices $u, v \in N_1$ and that there exists a vertex $b \in N_2$ adjacent neither to u nor to v. As before, let x' be a neighbor of x outside C. We have a satellite around $bwuvxx'$ with a path of length 3 on it, say $x'xuw$, a contradiction. □

Now we are able to characterize those graphs that are not biconnected and admit the empty SP-sufficient set:

Theorem 1 *Let G be an undirected graph that is connected but not biconnected. The empty set is SP-sufficient for G if and only if the following hold:*

(i) The block graph of G is a chain.

(ii) Every block of G that contains two cut-vertices is a clique.

(iii) Every block C of G that contains only one cut-vertex x does not contain a C^4, N_1^C and N_2^C are cliques and if there exists a vertex $w \in N_2^C$ adjacent to two vertices $u, v \in N_1^C$, then all vertices in N_2^C are adjacent to u or to v, where N_1^C, N_2^C are the subgraphs of G induced by the neighborhoods of x in C in distances 1, 2 respectively. Moreover, no vertex in C is in distance 3 from x.

Proof. The "only if" part is clear from the previous lemmas. For the "if" part we show how to color a set $\mathcal{P}$ of shortest paths on a given graph G that satisfies the conditions of the statement with $L(\mathcal{P})$ colors. A high level description of our approach follows. We will first modify $\mathcal{P}$ by shortening some paths and discarding some others. We will then construct a directed graph G' and a set $\mathcal{P}'$ of directed paths on G' in 1-1 correspondence with $\mathcal{P}$, with $L(\mathcal{P}') = L(\mathcal{P})$ and with the same conflict graph. We will obtain a coloring for $\mathcal{P}$ with $L(\mathcal{P})$ colors by coloring $\mathcal{P}'$ with $L(\mathcal{P}')$ colors. The coloring of $\mathcal{P}'$ will be done by computing many different local colorings that will be merged to give a single global coloring for $\mathcal{P}'$. A coloring for the initial set of paths will be obtained by coloring greedily the paths that were discarded in the first phase.

Let $C_1, \dots, C_k$ be the blocks of G ordered so that C_1, C_k are the two blocks containing only one cut-vertex and $|C_i \cap C_{i+1}| = 1$ for all $1 \le i < k$. Let $x_1, \dots, x_{k-1}$ be the cut-vertices of G ordered so that $C_i \cap C_{i+1} = \{x_i\}$. Let $\mathcal{P}^1 \subseteq \mathcal{P}$ be the set of all paths in $\mathcal{P}$ of length one. We remove these paths from $\mathcal{P}$. Now consider all paths that use an edge in one of $N_2^{C_1}, N_2^{C_k}$. Since we are dealing with shortest paths no such path can use an edge incident to x_1 or x_{k-1}. Therefore, all these paths are of length at most two. We claim that every edge in $N_2^{C_1}, N_2^{C_k}$ will only be used by identical paths. To see this assume that there is an edge uv in, say $N_2^{C_1}$ (the case for C_k is similar), that is used by two paths that are non-identical. Since $N_2^{C_1}$ is a clique and these paths are shortest they must go from u or v to two different vertices of $N_1^{C_1}$, say w, w'. There are two cases: either these paths are wuv and $w'vu$ or wuv and $w'uv$. Since these paths are shortest, in the first case C_1 must contain a C^4, while in the second case $u \in N_2^{C_1}$ is adjacent to both $w, w' \in N_1^{C_1}$ and there exists a vertex in $N_2^{C_1}$, namely v, that is not adjacent to any of w, w'. Both cases contradict the third condition of the statement and hence we obtain that every edge in $N_2^{C_1}, N_2^{C_k}$ will only be used by identical paths of length 2. Hence, we can shorten all such paths so that they do not use any edge in $N_2^{C_1}, N_2^{C_k}$ without modifying the conflict graph of $\mathcal{P}$. If, after this modification, we obtain any paths of length one we remove them from $\mathcal{P}$ and add them to $\mathcal{P}^1$.

Now we proceed to show how we construct G' and the set $\mathcal{P}'$ of directed paths on G' with the properties described above. An example is shown in Fig. 3. For each block C_i of G containing two cut-vertices x_{i-1}, x_i we construct its corresponding gadget as follows. We take two copies v_{in}, v_{out} of every vertex v of C_i that is not a cut-vertex and connect x_{i-1} with a directed edge to each *in* vertex and each *out* vertex with a directed edge to x_i. Finally, we connect x_{i-1} with a directed edge to x_i. The gadget for the blocks that contain only one cut-vertex, say for C_1, is constructed by starting from C_1, deleting all edges within each of $N_1^{C_1}, N_2^{C_1}$ and orienting the rest of the edges from $N_2^{C_1}$ to $N_1^{C_1}$ and from $N_1^{C_1}$ to x_1. Finally we add edges from each vertex of $N_1^{C_1}$ to all other vertices of $N_1^{C_1}$. The gadget for C_k is constructed similarly as for C_1 but with reverse orientations for the edges. To complete the construction of G' we connect all gadgets by identifying those vertices that correspond to the same cut-vertices in G. Since we have removed all paths of length one and have modified $\mathcal{P}$ so that no path uses an edge in $N_2^{C_1}$ and $N_2^{C_k}$, every path in $\mathcal{P}$ has a unique correspondent in G' and hence the construction of the set of paths $\mathcal{P}'$ on G' is straightforward. Notice that every edge in $N_1^{C_1}$ and $N_1^{C_k}$ corresponds to two oppositely directed edges in G'. However, since C_1 and C_k contain no C^4, all paths that use an edge e in $N_1^{C_1}$ or $N_1^{C_k}$ correspond to paths in G' that all use the edge, corresponding to e, in the same direction and hence the conflict graph does not change.

We continue with the coloring of $\mathcal{P}'$. Let $y_1, \ldots, y_l$ be the vertices in G' that correspond to vertices in $N_1^{C_1}$ and let $z_1, \ldots, z_m$ be the vertices in G' that correspond to vertices in $N_1^{C_k}$. Define $V' := \{x_1, \ldots, x_{k-1}, y_1 \ldots, y_l, z_1, \ldots, z_m\}$. For every vertex $v \in V'$, let Q_v^{out} be the set of edges directed out of v and Q_v^{in} be the set of edges directed into v. For $1 \le i < k$ let $\mathcal{P}'_{x_i}$ be the set of paths that touch x_i (i.e., start at x_i, end at x_i or go over x_i), for $1 \le i \le l$ let $\mathcal{P}'_{y_i}$ be the set of paths that go over y_i or start at y_i, and for $1 \le i \le m$ let $\mathcal{P}'_{z_i}$ be the set of paths that go over z_i or end at z_i. Notice that $\cup_{v \in V'} \mathcal{P}'_v = \mathcal{P}'$. We show how to color $\mathcal{P}'_v$ for all $v \in V'$ with $L(\mathcal{P}'_v)$ colors. Our method is similar to the one given in [3]. We build a bipartite multigraph H_v with one vertex for each edge incident to v. The one side of the partition consists of the vertices corresponding to edges in Q_v^{out} and the other consists of the vertices corresponding to edges in Q_v^{in}. Each path in $\mathcal{P}'_v$ that goes over v uses one edge in Q_v^{in}, one in Q_v^{out} and no other edge incident to v. For each such path we add one edge in H_v connecting its *out* vertex to its *in* vertex. For every path that starts or ends at v we add a loop to the vertex of H_v that corresponds to the edge incident to v used by this path. By König's classical result [8], H_v has a proper edge-coloring with $\Delta(H_v)$ colors which

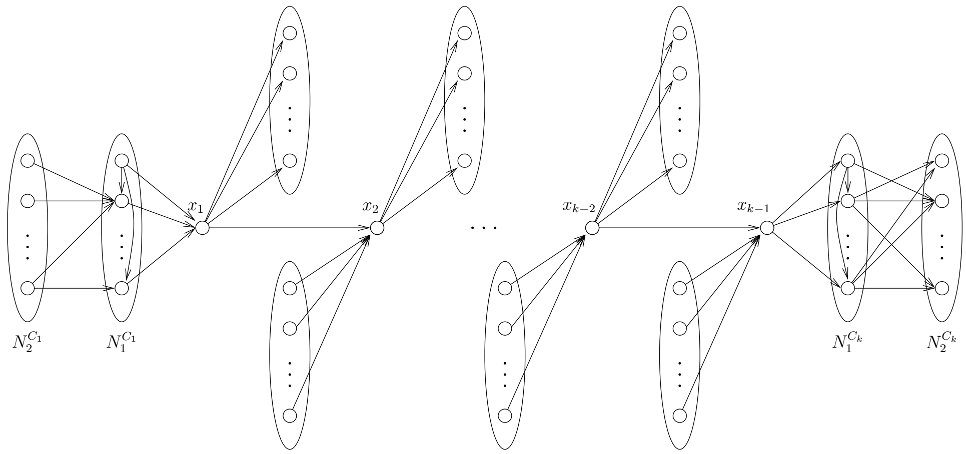

Figure 3: Example of the construction used in the proof of Theorem 1.

we can find in polynomial time. Since two paths in $\mathcal{P}'_v$ intersect in G' if and only if they intersect in an edge incident to v this results in a valid coloring of $\mathcal{P}'_v$. Furthermore, this coloring uses $L(\mathcal{P}'_v)$ colors since $\Delta(H_v) = L(\mathcal{P}'_v)$ (we assume that a loop contributes 1 to the degree of the vertex).

Let S_v be the coloring obtained for $\mathcal{P}'_v$. We show how we can merge all local colorings S_v into a global coloring S for $\mathcal{P}'$ without increasing the maximum number of colors used. We start by merging S_{x_1} with each of S_v for $v \in \{y_1, \ldots, y_l\}$. The merging for some S_{y_i} is done as follows. The only paths that are in both S_{x_1} and S_{y_i}, and thus might cause conflicts when merging, are the paths that use edge (y_i, x_1). Notice that if a path p colored in S_{y_i} intersects a path q that is colored in S_{x_1} but not in S_{y_i}, then p uses (y_i, x_1) and both p, q use the same outgoing edge incident to x_1, and are both colored with different colors in S_{x_1}. Therefore, to combine the two colorings we maintain S_{x_1} as is and modify S_{y_i}: we permute S_{y_i} (i.e., we rename the colors) so that the paths that use edge (y_i, x_1) get the color they have in S_{x_1}. Since S_{y_i} was a valid coloring, it remains valid after the permutation. Now the two colorings are compatible and can be merged in the obvious way. After the merging, S_{x_1} is the new coloring. After l mergings we have extended S_{x_1} to include the colorings $S_{y_1}, \ldots, S_{y_l}$ without increasing the maximum number of colors used, while maintaining its validity. We repeat the same procedure to merge coloring $S_{x_{k-1}}$ with the colorings $S_{z_1}, \ldots, S_{z_m}$.

Now, we can merge $S_{x_1}, \ldots, S_{x_{k-1}}$ to one global coloring S for $\mathcal{P}'$. To do this we initially set $S = S_{x_1}$ and continue to the processing of S_{x_2}. After processing S_{x_i} we continue to $S_{x_{i+1}}$ and merge the previous global coloring S (which we have obtained by merging colorings $S_{x_1} \ldots S_{x_i}$) with $S_{x_{i+1}}$. The merging of S with $S_{x_{i+1}}$ is done as follows. Consider the paths that use edge (x_i, x_{i+1}) on G'. These paths are the only paths that are in both S and $S_{x_{i+1}}$ and therefore might cause conflicts in the merging. To combine the two colorings we maintain S as is and modify $S_{x_{i+1}}$: we permute $S_{x_{i+1}}$ so that the paths that use edge (x_i, x_{i+1}) get the color they have in S. Since $S_{x_{i+1}}$ was a valid coloring before the modification, it remains valid and now we can merge the two colorings in the obvious way. This way we extend coloring S without increasing the number of colors used in S and $S_{x_{i+1}}$ while maintaining its validity. After $k-2$ mergings we obtain a global coloring S for $\mathcal{P}'$ that uses $\max_{1 \le i \le k-1} L(\mathcal{P}'_i) = L(\mathcal{P}')$ colors. We obtain a coloring for the initial set $\mathcal{P}$ with $L(\mathcal{P})$ colors by extending S greedily to the paths in $\mathcal{P}^1$ which we have previously discarded. $\square$

4 Biconnected Graphs

Now, consider the case where G is biconnected. We assume that G admits the empty SP-sufficient set.

Lemma 6 *For all $u, v \in V$: $d(u, v) \leq 3$.*

Proof. Assume to the contrary that $d(u, v) > 3$ and let $ua_1a_2 \ldots a_kv$ and $ub_1b_2 \ldots b_lv$ be two disjoint $u - v$ paths, each of length at least 4, such that $k + l$ is minimum. Notice that if there are edges $a_ib_{j'}, a_{i'}b_j$ with $i < i'$ and $j < j'$ then $i' = i + 1$ and $j' = j + 1$ since otherwise $k + l$ is not minimum. Also, if there are two such edges then we also have $a_ib_j, a_{i+1}b_{j+1} \in E$ since if this were not the case we would have a claw. We obtain a contradiction by case analysis depending on the edges between a_1, b_2 and a_2, b_1, as in the proof of Lemma 2(i). There are four main cases:

Case 1: *$a_1b_2, b_1a_2 \in E$.*
We have that $a_1b_1, a_2b_2 \in E$. Notice that $a_1b_3, b_1a_3 \notin E$ since otherwise $k+l$ would not be minimum. Therefore, we have that $d(u, b_3) = 3$ since by our choice of the two $u - v$ paths the only shortcut between u and b_3 can be through a_1. We distinguish cases depending on the edges between a_2, b_3 and b_2, a_3, as shown in Fig. 4(a).

Case 1.1: *$a_2b_3, b_2a_3 \notin E$.*
We have a net around $ub_1a_2b_2b_3a_3$ or an antenna around the same vertices if $a_3b_3 \in E$. Moreover, we can find a path of length 3 on the net or the antenna, namely the path $ub_1b_2b_3$ (recall that $d(u, b_3) = 3$), a contradiction.

Case 1.2: *$a_2b_3 \in E, b_2a_3 \notin E$.*
We have a satellite around $ub_1a_2b_2b_3b_4$ and a path of length 3 on it, namely the path $ub_1b_2b_3$. If $b_4 = v$ (i.e., if $l = 3$), then the satellite can not be broken and we obtain a contradiction. Otherwise we should have $a_2b_4 \in E$ but this creates a tent around $ub_1a_2b_2b_3b_4b_5$ and if $b_5 = v$ we have a contradiction. Otherwise $a_2b_5 \in E$ and we obtain a contradiction due to a claw at a_2.

Case 1.3: *$a_2b_3 \notin E, b_2a_3 \in E$.*
Similar to case 1.2.

Case 1.4: *$a_2b_3, b_2a_3 \in E$.*
We have $a_3b_3 \in E$ and we get a satellite around $ub_1a_2b_2b_3b_4$ with a path of length 3 on it, namely the path $ub_1b_2b_3$. This gives us a contradiction since $a_2b_4 \notin E$ because otherwise $k + l$ would not be minimum and hence the satellite can not be broken.

Case 2: *$a_1b_2 \in E, b_1a_2 \notin E$.*
We have $a_2b_2 \in E$ because of a claw at a_1. Also, notice that we can not have $b_1a_3 \in E$ since in that case $k + l$ would not be minimum. By our choice of the two $u - v$ paths the only shortcut from u to a_3 could be through b_1. It follows that $d(u, a_3) = 3$. We distinguish 4 cases depending on the edges between a_2, b_3 and b_2, a_3, as shown in Fig. 4(b).

Case 2.1: *$a_2b_3, b_2a_3 \in E$.*
We have a satellite around $ua_1b_2a_2a_3a_4$. Moreover there exists a path of length 3 on it, namely $ua_1a_2a_3$, since $d(u, a_3) = 3$ as was shown above. Hence, if $a_4 = v$ we have a contradiction. Otherwise we must have $b_2a_4 \in E$ but this can not be the case since then $k+l$ is not minimum, a contradiction.

Case 2.2: *$a_2b_3 \in E, b_2a_3 \notin E$.*
We have a satellite around $ua_1b_2a_2b_3b_4$. If $b_4 = v$ the satellite can not be broken and we can find

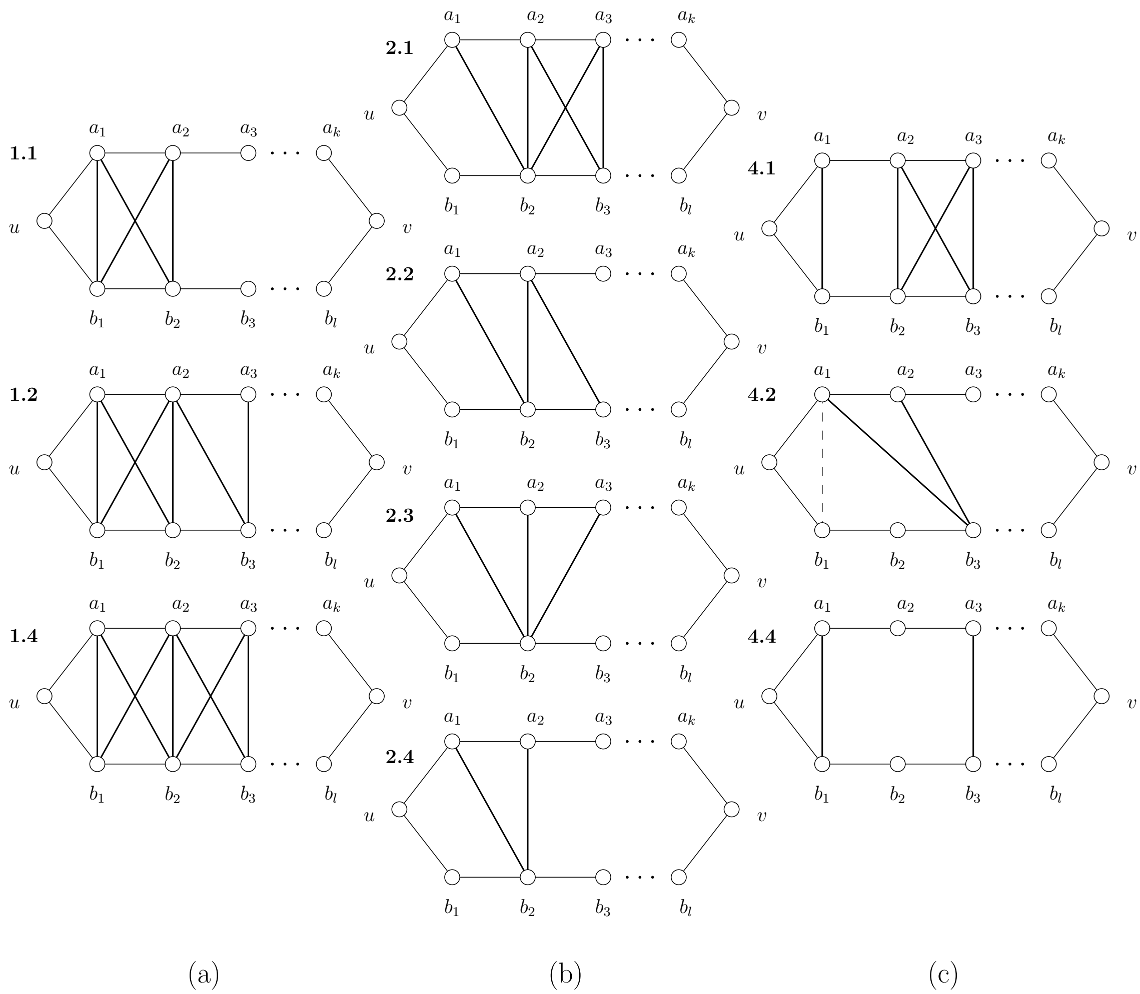

Figure 4: Case analysis in the proof of Lemma 6.

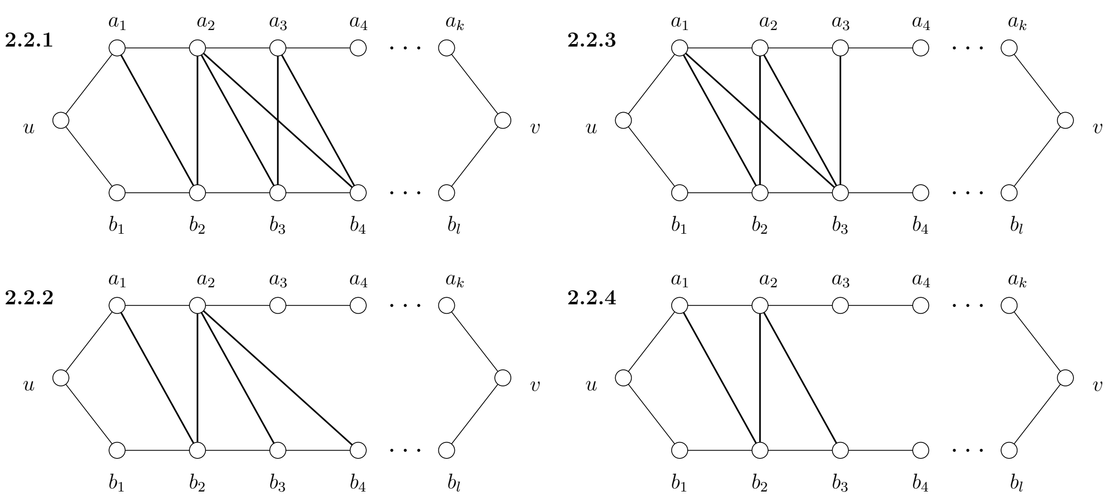

Figure 5: Case analysis (case 2.2) in the proof of Lemma 6.

the path $ua_1b_2b_3$ of length 3 on it, since $d(u,v) \geq 4$ and hence $a_1b_3 \notin E$. So we can assume $l \geq 4$. We distinguish more cases, according to the edges that are present between a_2, b_4 and a_3, b_3. We have 4 cases shown in Fig. 5.

Case 2.2.1: $a_2b_4, a_3b_3 \in E$.

We have $a_3b_4 \in E$ and there is a satellite around $b_1b_2a_2b_3b_4b_5$. Moreover there is a path of length 3 on it, namely the path $b_1b_2b_3b_4$ since $a_1b_4 \notin E$ and $b_1a_i \notin E$ for all $i \geq 2$. Hence, if $b_5 = v$ the satellite can not be broken and we obtain a contradiction. Otherwise, we should have $a_2b_5 \in E$ but this gives a tent around $b_1b_2a_2b_3b_4b_5b_6$ and a contradiction in the case where $b_6 = v$. If $b_6 \neq v$ then we should have $a_2b_6 \in E$ but this gives a claw at a_2, a contradiction.

Case 2.2.2: $a_2b_4 \in E, a_3b_3 \notin E$.

Similar to case 2.2.1.

Case 2.2.3: $a_2b_4 \notin E, a_3b_3 \in E$.

We have that $a_1b_3 \in E$ due to a satellite around $ua_1b_2a_2b_3b_4$. This creates another satellite around $ua_1b_3a_2a_3a_4$ with a path of length 3 on it, namely the path $ua_1a_2a_3$. If $a_4 = v$ this gives us a contradiction; otherwise we should have $b_3a_4 \in E$. This creates a tent around $ua_1b_3a_2a_3a_4a_5$. If $a_5 = v$ we have reached a contradiction; otherwise we should have $a_5b_3 \in E$ but this creates a claw at b_3, a contradiction.

Case 2.2.4: $a_2b_4, a_3b_3 \notin E$.

We have a net around $b_1b_2a_2a_3b_3b_4$ or an antenna around the same vertices if $a_3b_4 \in E$. Moreover there is a path of length 3 on the net or the antenna, namely the path $b_1b_2b_3b_4$ since $a_1b_4 \notin E$ and $b_1a_i \notin E$ for all $i \geq 2$. Hence, a contradiction.

Case 2.3: $a_2b_3 \notin E, b_2a_3 \in E$.

We have a claw at b_2 that can not be broken, a contradiction.

Case 2.4: $a_2b_3, b_2a_3 \notin E$.

We have a claw at b_2 that can not be broken, a contradiction.

Case 3: $a_1b_2 \notin E, b_1a_2 \in E$.

Similar to case 2.

Case 4: $a_1b_2, b_1a_2 \notin E$.

Again, we distinguish four cases according to the edges between a_2, b_3 and b_2, a_3. Three of these cases are shown in Fig. 4(c).

Case 4.1: $a_2b_3, b_2a_3 \in E$.

We have that $a_2b_2, a_3b_3 \in E$. Because of the C^5 around $ub_1b_2a_2a_1$ we should have $a_1b_1 \in E$. We have an antenna around $a_1b_1b_2a_2a_3a_4$ with a path of length 3 on it, namely the path $b_1b_2a_3a_4$ since $b_1a_5, b_1a_4, b_1a_3, b_2a_4 \notin E$ because otherwise $k+l$ would not be minimum. Moreover the antenna can not be broken and we reach a contradiction.

Case 4.2: $a_2b_3 \in E, b_2a_3 \notin E$.

We distinguish two cases according to whether $a_1b_1 \in E$.

Case 4.2.1: $a_1b_1 \in E$.

We have a C^5 around $a_1a_2b_3b_2b_1$. Thus, $a_2b_2 \in E$ or $a_1b_3 \in E$. Notice that if $a_2b_2 \in E$ we get a claw at a_2 that can not be broken. Hence, $a_1b_3 \in E$ and we get claws at b_3 that force $a_1b_4, a_2b_4 \in E$ (and hence we must have $l \geq 5$; otherwise $d(u,v) < 4$). This creates an antenna around $a_1b_1b_2b_3b_4b_5$ with a path of length 3 on it, namely the path $b_5b_4a_1b_1$, since $a_1b_5, b_1a_j \notin E$ for all $j \geq 2$ because otherwise $k+l$ would not be minimum. Furthermore, the antenna can not be broken and thus we reach a contradiction.

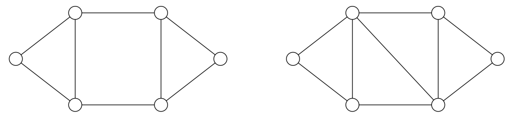

Figure 6: Examples of graphs of diameter 3 that admit the empty SP-sufficient set.

Case 4.2.2: $a_1b_1 \notin E$.
We have a C^6 around $ua_1a_2b_3b_2b_1$ and if $a_1b_3 \notin E$ we can find a path of length 3 on it, namely the path $ub_1b_2b_3$. Note that $a_2b_2 \notin E$ due to a claw at a_2. Hence, $a_1b_3 \in E$. Then we have a C^5 around $ua_1b_3b_2b_1$ that can not be broken, a contradiction.

Case 4.3: $a_2b_3 \notin E, a_3b_2 \in E$.
Similar to case 4.2.

Case 4.4: $a_2b_3, b_2a_3 \notin E$.
We have that $a_1b_3, b_1a_3 \notin E$ since otherwise we would have a claw at a_1, b_1 respectively. Hence, the paths of length 3 $ub_1b_2b_3$ and $ua_1a_2a_3$ are shortest. Therefore, if $a_1b_1 \notin E$ we have a C^5 or a longer induced cycle with a path of length 3 on it. It follows that $a_1b_1 \in E$. We claim that the paths $b_1b_2b_3b_4$ and $a_1a_2a_3a_4$ are also shortest. We will show it for $b_1b_2b_3b_4$ (it follows with similar arguments for $a_1a_2a_3a_4$). Recall that $b_1a_3, a_1b_3 \notin E$. We have that $a_1b_4 \notin E$ because otherwise we have a C^5 around $a_1b_1b_2b_3b_4$. Also, if $b_1a_4 \in E$ we have a C^5 around $a_1a_2a_3a_4b_1$. Thus, if there exists a shortcut from b_1 to b_4 it should be through a_j with $j \geq 5$. However, in this case we would have a C^6 or greater with a path of length 3 on it, namely $a_2a_3a_4a_5$ since there could not be another shortcut from a_2 to a_5 through some b_i because then $k + l$ would not be minimum. Hence, the paths $b_1b_2b_3b_4, a_1a_2a_3a_4$ are shortest. Thus, we should have $a_2b_2 \in E$ or $a_3b_3 \in E$ since otherwise we would have a C^5 or a longer induced cycle with a path of length 3 on it. If $a_2b_2 \in E$ we have claws at a_2 and b_2. Hence, $a_3b_3 \in E$ and this gives a C^6 around $a_1a_2a_3b_3b_2b_1$. We reach a contradiction since we can find a path of length 3 on the cycle between a_1 and b_3.

We considered all possible cases, reaching a contradiction for each of them; the lemma follows. □

Notice that the bound in the statement is tight: there exist graphs that admit the empty SP-sufficient set and have diameter 3. Two such graphs are shown in Fig. 6. On the other hand, if the diameter is at most two (or if we restrict to shortest-path routings of length at most two), then we can efficiently check whether a graph admits the empty SP-sufficient set as the following theorem illustrates.

Theorem 2 *Let $G = (V, E)$ be an undirected graph of diameter at most 2. The empty set is SP-sufficient for G if and only if E can be 2-colored such that every shortest path of length 2 in G uses edges of different colors.*

Proof. For sufficiency we show how to color a set $\mathcal{P}$ of shortest paths on G with $L(\mathcal{P})$ colors. We assume that $\mathcal{P}$ contains only paths of length two since any paths of length one can be colored greedily. We construct a multigraph H on E and for each path uvw in $\mathcal{P}$ we add an edge in H between uv and vw. If the condition holds then H is bipartite and by König's theorem [8] it can be edge-colored with $\Delta(H)$ colors in polynomial time. Since the degree of each vertex of H is equal to the load of the corresponding edge of G this gives a coloring for $\mathcal{P}$ with $L(\mathcal{P})$ colors.

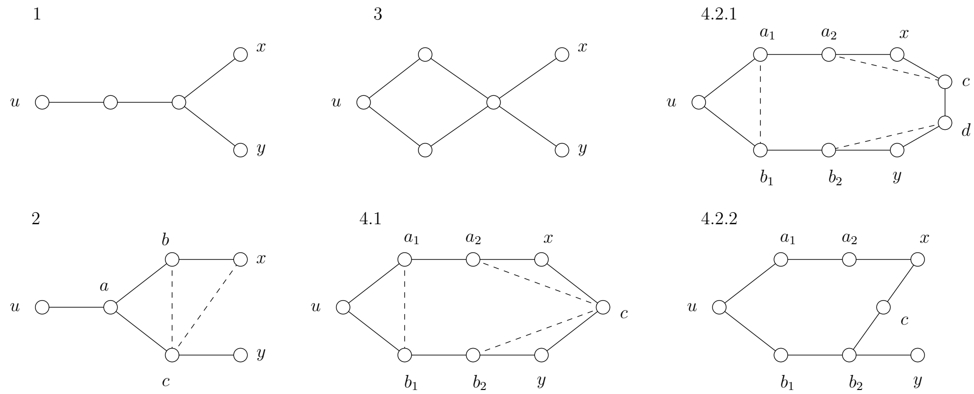

Figure 7: Case analysis in the proof of Lemma 7.

To see that the condition is necessary assume that it does not hold and let $\mathcal{P}$ be the set of all shortest paths of length two on G. We construct a multigraph H on E as before. Since E can not be 2-colored such that every path in $\mathcal{P}$ uses edges of different colors, H is not bipartite and hence contains an odd cycle. The paths corresponding to the edges of the odd cycle form a set of paths $\mathcal{P}'$ with load $L(\mathcal{P}') = 2$ and $\chi(\mathcal{P}') = 3$. □

With this result at our disposal, we can assume from now on that $\mathrm{diam}(G) = 3$. For $u \in V$, we denote with $N_i(u)$ the subgraph of G induced by the neighborhood of u in distance i. With the following 5 lemmas we establish that, for a vertex u of eccentricity 3, $N_1(u), G[V(N_2(u)) \cup V(N_3(u))]$ are cliques.

Lemma 7 *If u has eccentricity 3, then $N_3(u)$ is a clique.*

Proof. Let u be a vertex of eccentricity 3, x, y two vertices in distance 3 from u, and assume to the contrary that $xy \notin E$. Let p, q be two paths of length 3 connecting u to x and y respectively. We distinguish cases according to the number of common edges and vertices between p and q. The different cases are shown in Fig. 7.

Case 1: *p and q share 2 edges.*
We must have $xy \in E$ because of a claw.

Case 2: *p and q share one edge.*
Let $uabx, uacy$ be the two paths. We have a claw at a and hence $bc \in E$. This creates a net and so we should have $cx \in E$ or $by \in E$. Assume w.l.o.g. $cx \in E$. We get a claw at c which can not be broken, a contradiction.

Case 3: *p and q share no edges but have a common vertex.*
Again, we should have $xy \in E$ because of a claw.

Case 4: *p and q share no edges and no vertices.*
Let a_1, a_2 be the internal vertices of p and b_1, b_2 the internal vertices of q. If we have any of the edges $a_1b_2, b_1a_2, a_2y, b_2x$, then we can apply one of the cases 1-3. So we can assume now that none of these edges is present. Also, we can not have $a_2b_2 \in E$ either, since that would create a claw. We distinguish two cases depending on $d(x, y)$.

Case 4.1: $d(x, y) = 2$.
We have to consider only the case where there is a path of length two between x, y, disjoint from p and q, since otherwise we can apply case 1. Let c be the internal vertex of such a path. If $a_1c \in E$ or $b_1c \in E$ we can apply case 1 and thus we can assume that $a_1c, b_1c \notin E$. Hence, the only possible edges are a_1b_1, a_2c, b_2c. Note also that $d(u, c) = 3$ because otherwise we can apply case 1. It is easy to see that in all possible cases there is an induced odd cycle or an induced even cycle with a path of length 3 on it. The only case where there is an induced even cycle with no obvious path of length 3 on it, is when $a_1b_1, a_2c \in E$ (the case where $a_1b_1, b_2c \in E$ is identical). However, for this case we have that $d(a_1, y) = 3$ since otherwise we can apply case 2.

Case 4.2: $d(x, y) = 3$.
We consider two different cases depending on whether there is an $x - y$ path, disjoint from p and q.

Case 4.2.1: *there is a path $xcdy$ disjoint from p and q.*
If $a_1d \in E$ or $b_1c \in E$ then we can apply case 2 and thus we can assume that $a_1d, b_1c \notin E$. If $a_1c \in E$ or $b_1d \in E$ then we get claws at c, d respectively. If $a_2d \in E$ or $b_2c \in E$ then we get claws at a_2, b_2 respectively. Thus, the only possibilities are a_1b_1, a_2c, b_2d. In all cases we have an induced odd cycle or an induced even cycle with a path of length 3 on it. The only exception is the case where all three edges are there, i.e., $a_1b_1, a_2c, b_2d \in E$. For this case we show that $d(a_2, b_2) = 3$. Assume to the contrary that there exists an $a_2 - b_2$ path of length 2. Since we can not have any other edges between existing vertices, this path uses a new vertex, say w. We have a C^5 around $a_1a_2wb_2b_1$ and therefore we should have $a_1w \in E$ or $b_1w \in E$. Assume w.l.o.g. that $a_1w \in E$. If $wy \in E$ we can apply case 2 and thus we can assume that $wy \notin E$. Due to a claw at b_2 we get $b_1w \in E$. Due to the C^5 around a_2wb_2dc we should have $wc \in E$ or $wd \in E$. In both cases we get a claw at w with no possibility to break it, a contradiction.

Case 4.2.2: *there is no $x - y$ path of length 3 disjoint from p and q.*
There is a path of length 3 that has one common edge with p or q, say xcb_2y. We have that $cy \notin E$ since $d(x, y) = 3$. Also $b_1c \notin E$ since otherwise we can apply case 2. Therefore we have a claw at b_2 that can not be broken, a contradiction.

We considered all possible cases, reaching a contradiction for each of them; the lemma follows. □

Lemma 8 *If u has eccentricity 3, then $N_2(u)$ is a clique.*

Proof. Let u be a vertex of eccentricity 3, v a vertex at distance 3 from u, and x, y two vertices in distance 2 from u. Assume to the contrary that $xy \notin E$. There exist two disjoint paths of length 2 from u to x and y; otherwise we have a claw that forces $xy \in E$. Let uax, uby be two such paths. We distinguish cases depending on the distance of x, y from v. From Lemma 6 we have that their distance from v is at most 3. Notice that if $d(x, v) = d(y, v) = 3$ then, since by Lemma 7 $N_3(v)$ is a clique, we have $xy \in E$, a contradiction. Some of the cases are shown in Fig. 8. Notice that in all cases $ay, bx \notin E$ due to claws. Assume, w.l.o.g., $d(x, v) \leq d(y, v)$.

Case 1: $d(x, v) = d(y, v) = 1$.
Depending on whether $ab \in E$ we have a C^5 or a C^6 with a path of length 3 on it, a contradiction.

Case 2: $d(x, v) = 1, d(y, v) = 2$.
Let ycv be a path of length two connecting y and v. Notice that $ac \notin E$ because of a claw at c.

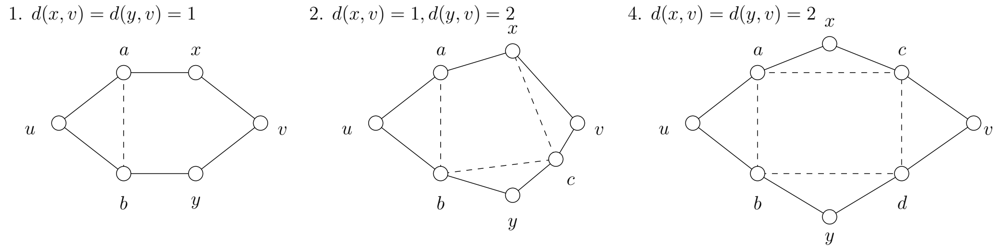

Figure 8: Case analysis in the proof of Lemma 8.

The only possible edges are ab, bc, xc. We distinguish cases according to which of these edges are present.

Case 2.1: *$ab, bc, xc \in E$.*
We can find a set of paths $\mathcal{P}$ with $L(\mathcal{P}) = 2$ and $\chi(\mathcal{P}) = 3$, namely the paths $ubcv$, vcy, ycx, cxa, xab, aby, ybu.

Case 2.2: *two of ab, bc, xc are present.*
We have a C^5.

Case 2.3: *one of ab, bc, xc is present.*
We consider the three cases.

Case 2.3.1: *$ab \in E$.*
We have a C^6 and we need to find a path of length 3 on it, in order to obtain a contradiction. We show that $d(b,v) = 3$. Assume for a contradiction that there exists a $b-v$ path of length 2. This path must use a new vertex, say w. Since $d(u,v) = 3$ we have that $uw \notin E$. Due to a claw at b we get $wy \in E$. This creates an antenna around $ubwycv$ with a path of length 3 on it, namely the path $ubwv$. In order to break it we should have $wc \in E$ since $uw \notin E$. We need $wa \in E$ or $wx \in E$ due to the C^5 around $abwvx$. In both cases we get a claw at w that can not be broken.

Case 2.3.2: *$bc \in E$.*
We have a C^6 with a path of length 3 on it, namely the path $ubcv$, a contradiction.

Case 2.3.3: *$xc \in E$.*
We have a C^6 and we need to find a path of length 3 on it, in order to obtain a contradiction. We show that $d(u,c) = 3$. Assume, as before, that there exists a $u-c$ path of length 2. This must use a new vertex, say w. Since $d(u,v) = 3$ we have $wv \notin E$. Because of a claw at c we get $wy \in E$. This creates an antenna around $vcywub$ with a path of length 3 on it, namely the path $vcwu$ and hence we should have $wb \in E$. Because of the C^5 around $uwcxa$ we should have $wa \in E$ or $wx \in E$. In both cases we get a claw at w that can not be broken, a contradiction.

Case 2.4: *none of ab, bc, xc is present.*
We have a C^7.

Case 3: $d(x,v) = 1, d(y,v) = 3$.
We have that $y, u \in N_3(v)$ and therefore by Lemma 7 we get $uy \in E$, a contradiction.

Case 4: $d(x,v) = d(y,v) = 2$.
If two paths of length two connecting x to v and y to v share a common edge then we have a claw. Thus, we can assume that there exist two disjoint paths xcv and ydv. We have that $ud, uc \notin E$ because $d(u,v) = 3$. If $xd, yc \in E$ we get claws at d and c respectively. If $ad, bc \in E$ we have again

claws at d, c. Thus, the only possibilities are ac, bd, ab, cd. We distinguish cases according to which of these edges exist.

Case 4.1: *$ac, bd, ab, cd \in E$.*
We can find a set of paths $\mathcal{P}$ with $L(\mathcal{P}) = 2$ and $\chi(\mathcal{P}) = 3$ similarly as in case 2.1.

Case 4.2: *three of ac, bd, ab, cd are present.*
We have a C^5.

Case 4.3: *two of ac, bd, ab, cd are present.*
We distinguish further cases.

Case 4.3.1: *$ac, bd \in E$.*
We get a C^6 with a path of length 3 on it, namely the path $uacv$.

Case 4.3.2: *$ab, ac \in E$.*
We take the following paths: $uacv$, cvd, vdy, dyb, yba, bax and xau.

Case 4.3.3: *$ab, bd \in E$ or $ac, cd \in E$ or $bd, cd \in E$.*
Similar to case 4.3.2.

Case 4.3.4: *$ab, cd \in E$.*
We have a C^6 around $abydcx$ and need to find a shortest path of length 3 on it in order to obtain a contradiction. If $d(x, y) = 3$ we are done. Otherwise there exists an $x - y$ path of length two in G. This path must use a new vertex, say w. Due to the C^5 around $axwyb$ we should have $aw \in E$ or $bw \in E$. Assume w.l.o.g. $bw \in E$. Due to the C^5 around $xcdyw$ we should have $wc \in E$ or $wd \in E$. If $wd \in E$ we get a claw at w. Hence, $wc \in E$. Also, $uw, vw, aw \notin E$ because of claws at w. We have two antennas around $vcwxab$ and $ubwydc$ and hence if $d(u, c) = 3$ or $d(v, a) = 3$ we are done. We will show that it can not be $d(v, a) = d(u, c) = 2$. Assume for a contradiction that there exist $v - a, u - c$ paths of length 2. The $v - a, u - c$ paths can not use any of the existing vertices. Let uz_1c, vz_2a be the two paths. We have a claw at a and $ux, uz_2 \notin E$, hence $xz_2 \in E$. This creates an antenna around $uaxz_2vc$ with a path of length 3 on it, namely the path uaz_2v and therefore we should have $z_2c \in E$. With similar arguments we obtain that $xz_1, z_1a \in E$. We can now construct a set of paths $\mathcal{P}$ with $L(\mathcal{P}) = 2$ and $\chi(\mathcal{P}) = 3$. If $z_1z_2 \notin E$ take uz_1x, z_1xz_2, xz_2v, vz_2au, uax, axc, xcv, vcz_1, and cz_1u. If $z_1z_2 \in E$ take uz_1z_2v, vz_2a, z_2au, uax, axc, xcv and vcz_1u.

Case 4.4: *one of ac, bd, ab, cd is present.*
We have a C^7.

Case 4.5: *$ac, bd, ab, cd \notin E$.*
In this case we have a C^8 and we need to find a shortest path of length 3 on it in order to obtain a contradiction. If $d(u, c) = 3$ we are done. Otherwise there exists a $u - c$ path of length two in G. This path must use a new vertex, say w. We have a claw at c and $wv \notin E$ since $d(u, v) = 3$. Therefore, $wx \in E$. We have an antenna around $vcxwua$ with a path of length 3 on it, namely the path $vcwu$ and so we should have $wa \in E$. We have a C^7 around $uwcvdyb$ but no possibility to break it because of claws at w, a contradiction. Thus, $d(u, c) = 3$ and we have a contradiction.

Case 5: $d(x, v) = 2, d(y, v) = 3$.
We have that $y, u \in N_3(v)$ and therefore by Lemma 7 we get $uy \in E$, a contradiction.

We considered all possible cases, reaching a contradiction for each of them; the lemma follows. □

Lemma 9 *Let $u, v \in V$ with $d(u, v) = 3$. There exist two disjoint $u - v$ paths, each of length 3.*

Proof. Since G is biconnected there exists a pair of disjoint $u - v$ paths. We will first show that there exists such a pair of paths such that at least one of them is of length 3. Assume to the

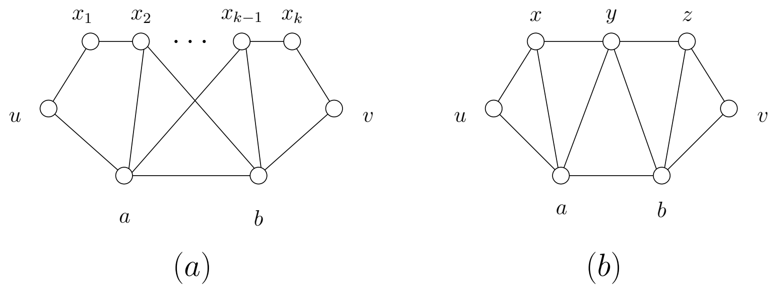

Figure 9: (a) Proof of Lemma 9. (b) The case $k = 3$.

contrary that every $u-v$ path of length 3 intersects all other $u-v$ paths in G. Let $uxyv$ be a $u-v$ path of length 3. Consider two disjoint $u-v$ paths $uxa_1 \ldots a_kv$ and $ub_1 \ldots b_lyv$, each of length at least 4 such that $k+l$ is minimum (since $uxyv$ crosses all pairs of disjoint $u-v$ paths and G is biconnected there exists such a pair of $u-v$ paths). Due to a claw at x and since $uy \notin E$ and $ua_1 \notin E$ we have $a_1y \in E$. Similarly, due to a claw at y we have $xb_l \in E$. Now we have a claw at x and since $a_1u \notin E$, $ub_l \notin E$ we get $a_1b_l \in E$. Therefore we have found two disjoint $u-v$ paths, namely $uxyv$ and $ub_1 \ldots b_la_1 \ldots a_kv$ such that one of them is of length 3, a contradiction.

We now proceed to show that there exists a pair of disjoint $u-v$ paths, each of length 3. Assume for a contradiction that this does not hold and let $uabv$ and $ux_1x_2 \ldots x_kv$ be two disjoint $u-v$ paths with $k \geq 3$ minimum, as shown in Fig. 9(a). Since $x_2, b \in N_2(u)$, $x_{k-1}, a \in N_2(v)$ and u, v are vertices of eccentricity 3, by Lemma 8 we have that $x_2b, x_{k-1}a \in E$. Because of claws in a and b we also have that $x_2a, bx_{k-1} \in E$. If $k \geq 5$ we have a claw at a that can not be broken, a contradiction. We consider the cases $k = 3$ and $k = 4$ separately. First assume $k = 4$. It can not be that $ax_4 \in E$ or $bx_1 \in E$ since that would create claws at a, b respectively. Therefore we can find a set of shortest paths $\mathcal{P}$ with $L(\mathcal{P}) = 2$ and $\chi(\mathcal{P}) = 3$: take $uabv$, uax_3, ax_3x_4, $x_4x_3x_2$, $x_3x_2x_1$, x_1x_2b, x_2bv. We continue for the case $k = 3$. Let $uabv, uxyzv$ be such a pair of disjoint $u-v$ paths as shown in Fig. 9(b). As before we have that $ya, yb \in E$. We have claws at y that force $xa \in E$ or $az \in E$ and $zb \in E$ or $xb \in E$. Notice that we can not have both $xb \in E$ and $za \in E$ since in that case we obtain two disjoint $u-v$ paths of length 3. Therefore, at least one of xa, zb should be present. Assume w.l.o.g. $xa \in E$. We have an antenna around $uaybvz$ with a path of length 3 on it, namely the path $uabv$ and hence we should have $bz \in E$ or $az \in E$. If $az \in E$ we get that $bz \in E$ because of a claw at a and hence we can assume that in every case we have $bz \in E$. If $xb, az \notin E$ then we are able to construct a set of shortest paths $\mathcal{P}$ with $L(\mathcal{P}) = 2$ and $\chi(\mathcal{P}) = 3$: take uab, abv, vby, byx, xyz, zya and yau. If one of xb, az is present, say $xb \in E$ (both cases are identical), we can still find such a set of paths provided $d(u, z) = 3$ (in the case where $az \in E$ we need $d(v, x) = 3$): $uxyz$, uxb, xbz, zba, bau, uay, ayz. It remains to show that $d(u, z) = 3$. Assume otherwise, i.e., $d(u, z) = 2$ (it can not be $uz \in E$ since $d(u, v) = 3$). Notice that since $az \notin E$ there must exist a vertex w with $uw, wz \in E$. But this gives us two disjoint $u-v$ paths both of length 3, a contradiction. □

Lemma 10 *If u has eccentricity 3, then $N_1(u)$ is a clique.*

Proof. Let u be a vertex of eccentricity 3, x, y be two neighbors of u, and v be such that $d(u, v) = 3$. Assume, for a contradiction, that $xy \notin E$. Clearly, $d(v, x) > 1$ and $d(v, y) > 1$. If $d(v, x) = d(v, y)$

then by Lemmas 7,8 we have $xy \in E$. The only remaining case is $d(v,x) = 2$ and $d(v,y) = 3$. By Lemma 9 we can assume that there exist two disjoint $y - v$ paths of length 3. Let ya_1a_2v and yb_1b_2v be two such paths (these paths can not use u since $d(u,v) = 3$). Since $x, a_1, b_1 \in N_2(v)$ we have $xa_1, xb_1, a_1b_1 \in E$. Also, $x, a_2, b_2 \in N_2(y)$ and thus $xa_2, xb_2, a_2b_2 \in E$. If $a_1b_2 \notin E$ and $b_1a_2 \notin E$, we can find a set of paths $\mathcal{P}$ with $L(\mathcal{P}) = 2$ and $\chi(\mathcal{P}) = 3$: take yb_1b_2v, vb_2x, b_2xa_1, xa_1y, ya_1a_2, a_1a_2v, va_2x, a_2xb_1, xb_1y. If $a_1b_2 \in E$ or $b_1a_2 \in E$ then the paths b_2xa_1 and b_1xa_2 are not shortest. In this case we can use instead of, say va_2x, a_2xb_1 and xb_1y, the path va_2b_1y. □

Lemma 11 *If u has eccentricity 3, then $G[V(N_2(u)) \cup V(N_3(u))]$ is a clique.*

Proof. Let u be a vertex of eccentricity 3 and x, y be such that $x \in N_2(u)$ and $y \in N_3(u)$. Assume that $xy \notin E$. By Lemma 9 we can assume that there exist two disjoint $u - y$ paths of length 3. Let ua_1a_2y and ub_1b_2y be two such paths. Since $x, a_2, b_2 \in N_2(u)$ we have $xa_2, xb_2, a_2b_2 \in E$. Also, $x, a_1, b_1 \in N_2(y)$ and thus $xa_1, xb_1, a_1b_1 \in E$. If $a_1b_2 \notin E$ and $b_1a_2 \notin E$, we can find a set of paths $\mathcal{P}$ with $L(\mathcal{P}) = 2$ and $\chi(\mathcal{P}) = 3$: take yb_2b_1u, ub_1x, b_1xa_2, xa_2y, ya_2a_1, a_2a_1u, ua_1x, a_1xb_2, xb_2y. If $a_1b_2 \in E$ or $b_1a_2 \in E$ then we can not use the paths b_1xa_2 and a_1xb_2. In this case we can use instead of, say ua_1x, a_1xb_2 and xb_2y, the path ua_1b_2y. □

For a vertex $v \in V$ of eccentricity 3 we define $T_G(v)$ to be the bipartite graph with vertex-sets $V_1 = V(N_1(v)) \setminus \{u \in N_1(v) \mid \mathrm{ecc}(u) = 3\}$, $V_2 = V(N_2(v))$ and edge-set $E_T = \{uw \in E \mid u \in V_1, w \in V_2\}$. We have the following theorem:

Theorem 3 *Let $G = (V, E)$ be an undirected, biconnected graph with* $\mathrm{diam}(G) = 3$ *and let $v \in V$ be a vertex of eccentricity 3. The empty set is SP-sufficient for G if and only if the following hold:*

(i) $N_1(v)$ and $G[V(N_2(v)) \cup V(N_3(v))]$ are cliques,

(ii) $T_G(v)$ contains no cycle, no P^5, and if it contains a P^4 then it has only one non-trivial connected component.

Proof. The necessity of (i) is clear from previous lemmas. To show that (ii) is also necessary we exhibit a set of paths $\mathcal{P}$ with load $L(\mathcal{P}) = 2$ and $\chi(\mathcal{P}) = 3$ for each case. First, assume that there is a cycle $a_1b_1a_2b_2 \dots a_kb_ka_1$ in $T_G(v)$ and assume w.l.o.g. $a_i \in V_1$ for $1 \leq i \leq k$. Let $u \in V$ be such that $d(u,v) = 3$. Consider the paths of length 3: $va_1b_1u, ub_1a_2v, \dots, va_kb_ku, ub_ka_1v$. This makes an even number of paths with load 2 whose conflict graph is a cycle; to obtain the required set of paths just replace one path with two paths of length two. Now, assume $T_G(v)$ contains a P^5 $a_1b_1a_2b_2a_3$, where $a_i \in V_1, 1 \leq i \leq 3$ as shown in Fig. 10(a). The paths va_1b_1u, ub_1a_2v, va_2b_2, a_2b_2u, ub_2a_3, $b_2a_3a_1$, $a_3a_1b_1$ have the desired property. Consider the case where $T_G(v)$ contains a P^4, say $a_1b_1a_2b_2$, and has more than one non-trivial connected components. There exists an edge xy in a different component than the one in which the P^4 is contained. Let $x, a_i \in V_1, 1 \leq i \leq 2$ as shown in Fig. 10(b). The paths va_2b_2u, va_2b_1u, b_2a_2x, a_2xy, yxa_1, xa_1b_1, a_1b_1u have the desired property.

For sufficiency we demonstrate how to color a set $\mathcal{P}$ of shortest paths on a graph G with $\mathrm{diam}(G) = 3$ that satisfies (i) and (ii) with $L(\mathcal{P})$ colors. We assume that $\mathcal{P}$ does not contain any paths of length one since these paths can be colored greedily after coloring the longer paths.

Define $Q_1 := \{u \in N_1(v) \mid \mathrm{ecc}(u) = 3\} \cup \{v\}$ and $Q_2 := V(N_3(v))$. Since $\mathrm{diam}(G) = 3$, we have that $Q_1 \cup Q_2 \cup V_1 \cup V_2 = V$, where V_1, V_2 are the vertex-sets of $T_G(v)$, as defined earlier (consider Fig. 11). Notice that by condition (i), $G[Q_1 \cup V_1]$ and $G[V_2 \cup Q_2]$ are cliques.

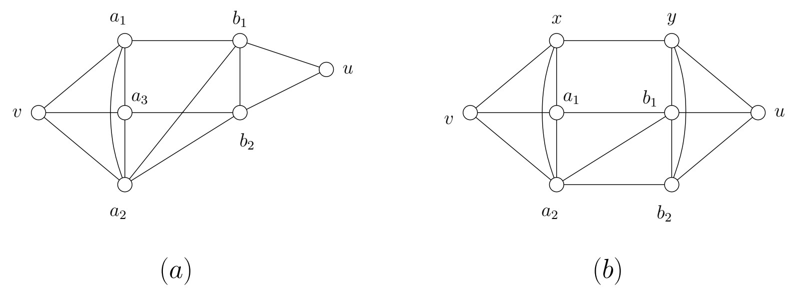

Figure 10: Proof of Theorem 3: $T_G(v)$ contains (a) a P^5, (b) a P^4.

We claim that all vertices of eccentricity 3 are contained in Q_1 and Q_2. By definition V_1 can not contain any vertices of eccentricity 3. Assume that there exists a vertex $u \in V_2$ with $\text{ecc}(u) = 3$. There exists a vertex $x \in V$ with $d(u, x) = 3$. Since $G[V_2 \cup Q_2]$ is a clique we have that $x \notin Q_2$. Also, since $d(u, v) = 2$, u has a neighbor in V_1 and since $G[Q_1 \cup V_1]$ is a clique we have that $x \notin Q_1 \cup V_1$. We reach a contradiction and thus V_2 does not contain any vertices of eccentricity 3.

First, assume that $T_G(v)$ does not contain a P^4. Let $C_1, \ldots, C_k$ be the connected components of $T_G(v)$ in some arbitrary ordering. Notice that at least one of $|V(C_i) \cap V_1| = 1$ or $|V(C_i) \cap V_2| = 1$ holds for all $1 \leq i \leq k$ (i.e., all components of $T_G(v)$ are stars), since no component of $T_G(v)$ contains a P^4 or a P^5. For each $1 \leq i \leq k$ let l_i be such that $|V(C_i) \cap V_{l_i}| = 1$ and x_i be such that $V(C_i) \cap V_{l_i} = \{x_i\}$. Let p be a path of length 3 in $\mathcal{P}$. Since Q_1, Q_2 contain all vertices of eccentricity 3 in G and both $G[Q_1], G[Q_2]$ are cliques, p goes from a vertex of Q_1 to a vertex of Q_2 and uses two edges incident to some x_i. Moreover, since all paths in $\mathcal{P}$ are shortest paths and $G[Q_1 \cup V_1], G[V_2 \cup Q_2]$ are cliques, a path in $\mathcal{P}$ intersects p in an edge not incident to x_i if and only if it intersects it in an edge incident to x_i. Therefore we can shorten p by restricting it to the edges incident to x_i without modifying the conflict graph of $\mathcal{P}$. We apply the same modification to all paths of length 3 in $\mathcal{P}$. Let $\mathcal{P}^1$ be the set of paths that after this modification have length one. These paths are discarded from $\mathcal{P}$. Now $\mathcal{P}$ contains only paths of length two.

In order to obtain a coloring of $\mathcal{P}$ with $L(\mathcal{P})$ colors it suffices to show that the edges that are used by paths in $\mathcal{P}$ can be 2-colored such that every path in $\mathcal{P}$ uses edges of different colors. Consider first the subgraph of G, $G[V_1 \cup Q_1]$, induced by $V_1 \cup Q_1$. No path in $\mathcal{P}$ uses two edges in this subgraph. Hence, we can color all these edges with one color, say a. The same holds for the subgraph $G[V_2 \cup Q_2]$. Moreover, since $\mathcal{P}$ contains no paths of length 3, no path touches edges from both $G[V_1 \cup Q_1]$, $G[V_2 \cup Q_2]$. Hence, we can color all the edges in $G[V_2 \cup Q_2]$ with color a. We color all remaining edges, i.e., edges in $T_G(v)$, with a second color b. No path uses two edges from $T_G(v)$ and therefore we can obtain a valid coloring of $\mathcal{P}$ with $L(\mathcal{P})$ colors using the reduction to bipartite edge-coloring illustrated in the proof of Theorem 2. We complete the coloring by coloring greedily the paths in $\mathcal{P}^1$ and thus obtain a valid coloring for $\mathcal{P}$ with $L(\mathcal{P})$ colors.

Now, assume that $T_G(v)$ contains a P^4. In this case, $T_G(v)$ consists of a single non-trivial connected component which is a double-star. Let $xyzw$ be a P^4 in $T_G(v)$. Let $\mathcal{P}_y, \mathcal{P}_z$ be the sets of paths on G that touch y, z respectively. Notice that $\mathcal{P}_y \cup \mathcal{P}_z = \mathcal{P}$. We color $\mathcal{P}$ in two steps: we first color $\mathcal{P}_y$ and $\mathcal{P}_z$ separately with $L(\mathcal{P}_y)$ and $L(\mathcal{P}_z)$ colors and then merge the two colorings into a global coloring for $\mathcal{P}$ without increasing the number of colors used. The coloring of the two sets of paths is done similarly for both $\mathcal{P}_y$ and $\mathcal{P}_z$. We demonstrate it for $\mathcal{P}_y$. Assume w.l.o.g.

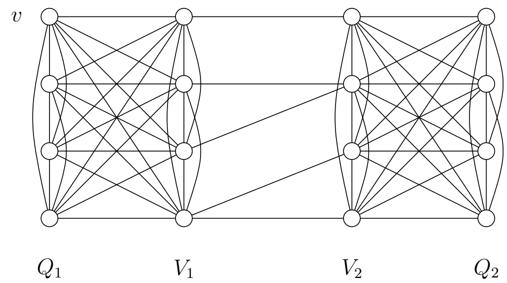

Figure 11: A biconnected graph of diameter 3 that does not require wavelength conversion.

$y \in V_2$. Two paths in $\mathcal{P}_y$ intersect in an edge not incident to y if and only if they intersect in an edge incident to y. Thus, we can restrict each path in $\mathcal{P}_y$ to the edges it uses that are incident to y. We modify each path this way and modify $\mathcal{P}_y$ accordingly. Let P_y^1 be the set of paths that after this modification are of length one. We discard these paths from $\mathcal{P}_y$. Hence, we need only exhibit a 2-coloring of the edges that are used by paths in $\mathcal{P}_y$ such that every path uses edges of different colors. Notice that every path in $\mathcal{P}_y$ uses either an edge connecting y to some vertex in Q_2 and an edge of $T_G(v)$ or an edge connecting y to some vertex in V_2 and an edge of $T_G(v)$. Therefore, we can color the edges of $T_G(v)$ with color a and the rest of the edges that are used by paths in $\mathcal{P}_y$ with color b. As before, we obtain a coloring of $\mathcal{P}_y$ with $L(\mathcal{P}_y)$ colors. The coloring is completed by coloring greedily the paths in $\mathcal{P}_y^1$.

In order to merge the two colorings into a valid global coloring we only need to take care of the coloring of the paths that are in both sets. The only paths with this property are paths that use edge yz. Also, a path in $\mathcal{P}_y$ intersects a path in $\mathcal{P}_z$ only if both paths use edge yz. Hence, if we permute the coloring of one set so that each path that is in both $\mathcal{P}_y$ and $\mathcal{P}_z$ gets the same color in both colorings, then the two colorings can be merged into a global valid coloring. This does not increase the number of colors used and therefore we obtain a coloring for $\mathcal{P}$ with $L(\mathcal{P})$ colors. □

Note that Theorem 3 along with Theorem 2 provide a complete characterization of the class of biconnected graphs that admit the empty SP-sufficient set since by Lemma 6 we know that all graphs with diameter greater than 3 need wavelength converters.

5 Conclusion and Future Work

We have given a complete characterization of undirected networks on which any set $\mathcal{P}$ of shortest paths admits a valid wavelength assignment with $L(\mathcal{P})$ wavelengths. These are exactly the networks that do not benefit from the use of wavelength conversion when shortest-path routing is used. It follows from our characterization that this class of networks is efficiently recognizable. We have also given an efficient algorithm for computing a wavelength assignment with $L(\mathcal{P})$ wavelengths for this class of networks. To our knowledge, these are the first theoretical investigations of a wavelength assignment problem in the practical scenario with shortest-path routing. The results should be contrasted with known results for arbitrary paths, because they suggest that the traditional worst-case analysis for arbitrary paths can yield overly pessimistic results concerning wavelength conversion requirements.

We note that a characterization of networks *with* converters that have the same property, i.e., admit a wavelength assignment with $L(\mathcal{P})$ wavelengths for any set $\mathcal{P}$ of shortest paths, does not follow directly from our current results. This is because the approach of Wilfong and Winkler [11] fails in our case: we can no longer "explode" the converters and argue about each component of the resulting graph independently, since such a modification will alter the distance between pairs of vertices. Therefore, it remains to see whether we can decide in polynomial time if a set is SP-sufficient for a given graph and also to determine the complexity of MINIMUM SP-SUFFICIENT SET and propose efficient exact or approximation algorithms for solving it. Also, a similar study should be carried out for bidirected or directed graphs.

References

[1] V. Auletta, I. Caragiannis, L. Gargano, C. Kaklamanis, and P. Persiano. Sparse and limited wavelength conversion in all-optical tree networks. *Theoretical Computer Science*, 266(1–2):887–934, 2001.

[2] T. Erlebach and K. Jansen. The complexity of path coloring and call scheduling. *Theoretical Computer Science*, 255(1–2):33–50, 2001.

[3] T. Erlebach, K. Jansen, C. Kaklamanis, M. Mihail, and P. Persiano. Optimal wavelength routing in directed fiber trees. *Theoretical Computer Science*, 221:119–137, 1999.

[4] T. Erlebach and S. Stefanakos. Wavelength conversion in networks of bounded treewidth. TIK-Report 132, ETH Zurich, April 2002. Available electronically at ftp://ftp.tik.ee.ethz.ch/pub/publications/TIK-Report132.pdf.

[5] M.R. Garey, D.S. Johnson, G.L. Miller, and C.H. Papadimitriou. The complexity of coloring circular arcs and chords. *SIAM Journal on Algebraic and Discrete Methods*, 1(2):216–227, 1980.

[6] L. Gargano. Limited wavelength conversion in all-optical tree networks. In *Proceedings of the 25th International Colloquium on Automata, Languages and Programming (ICALP '98)*, volume 1443 of *Lecture Notes in Computer Science*, pages 544–555, 1998.

[7] J.M. Kleinberg and A. Kumar. Wavelength conversion in optical networks. *Journal of Algorithms*, 38(1):25–50, 2001.

[8] D. König. Über Graphen und ihre Anwendung auf Determinantentheorie und Mengenlehre. *Math. Ann.*, 77:453–465, 1916.

[9] V. Kumar. Approximating circular arc colouring and bandwidth allocation in all-optical ring networks. In *Proceedings of the International Workshop on Approximation Algorithms for Combinatorial Optimization (APPROX '98)*, volume 1444 of *Lecture Notes in Computer Science*, pages 147–158, 1998.

[10] Y. Rabani. Path coloring on the mesh. In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science (FOCS'96)*, pages 400–409, 1996.

[11] G.T. Wilfong and P. Winkler. Ring routing and wavelength translation. In *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'98)*, pages 333–341, 1998.

[12] H. Zang, J.P. Jue, and B. Mukherjee. A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks. *Optical Networks Magazine*, 1(1):47–60, January 2000.