Detecting a Theta or a Prism

Maria Chudnovsky [1]
Columbia University, New York, NY 10027

Rohan Kapadia [2]
McGill University, Montreal, Canada

May 16, 2006; revised July 18, 2007

[1]This research was conducted during the period the author served as a Clay Mathematics Institute Research Fellow at Princeton University.
[2]This research was conducted at Princeton University

Abstract

A theta in a graph is an induced subgraph consisting of two nonadjacent vertices joined by three disjoint paths. A prism in a graph is an induced subgraph consisting of two disjoint triangles joined by three disjoint paths. This paper gives a polynomial-time algorithm to test whether a graph has an induced subgraph that is either a prism or a theta.

1 Introduction

The vertex set of a graph G is denoted by $V(G)$ and the edge set by $E(G)$. All graphs in this paper are simple. A *cycle* is a connected subgraph in which all vertices have degree two. A *path* P in G is an induced connected subgraph of G such that either P is a one-vertex graph, or two vertices of P have degree one and all others have degree two. (Please note that we are using a non-standard definition.) The *ends* of a path are the two vertices with degree one in the path. If $X \subseteq V(G)$, the subgraph with vertex set X and edge set all edges of G with both ends in X is denoted $G|X$, and called the subgraph *induced* on X.

A *prism* is a graph consisting of two disjoint triangles $\{a_1, a_2, a_3\}$ and $\{b_1, b_2, b_3\}$ together with three paths P_1, P_2 and P_3 such that the union of every two of P_1, P_2 and P_3 is an induced cycle. A *theta* is a graph consisting of two nonadjacent vertices s and t, and three paths P, Q and R with ends s and t, such that the union of every two of P, Q and R is an induced cycle. The vertices s, t are called the *ends* of the theta. Finally, a *pyramid* is a graph consisting of a triangle $\{b_1, b_2, b_3\}$, a vertex $a \notin \{b_1, b_2, b_3\}$, and, for $i = 1, 2, 3$ a path P_i between a and b_i, such that a is adjacent to at most one of b_1, b_2 and b_3 and the union of every two of P_1, P_2 and P_3 is an induced cycle. An induced subgraph H of G is a *prism, theta* or *pyramid in* G if H is a prism, theta or pyramid, respectively. We say that a graph G *contains* a prism, a theta or a pyramid, if some induced subgraph of G is a prism, a theta or a pyramid, respectively. A *hole* in a graph is an induced cycle of length at least four. A hole is *odd* if it has an odd number of vertices, and even otherwise. Given a graph G, the *complement* $\overline{G}$ of G is the graph with the same vertex set as G, and such that two vertices are adjacent in $\overline{G}$ if and only if they are non-adjacent in G. A graph G is called *Berge* if neither G nor $\overline{G}$ contains an odd hole.

In [1] Chudnovsky, Cornuéjols, Liu, Seymour, and Vušković present a polynomial time algorithm to test whether a graph contains a pyramid. The main result of [1] is a polynomial time algorithm to test if a graph is Berge. Since every pyramid contains an odd hole, no Berge graph contains a pyramid, and testing for pyramids is the first step in testing for Bergeness. Similarly to pyramids, every theta and every prism contains an even hole, and therefore a graph with no even holes contains no theta or prism. In [4] Maffray and Trotignon prove that the problem of deciding if a graph contains a prism is NP-complete, and in [2] Chudnovsky and Seymour show that testing if a graph contains a theta can be done in polynomial time. However, the complexity of deciding if a graph contains a prism or a theta remained open (and this is the analogue of the pyramid problem for graphs with no even holes). Our main result here is a polynomial time algorithm that given a graph G, decides if G contains either a prism or a theta. The algorithm is described in Section 5, and its running time is $O(|V(G)|^{35})$.

Before continuing with the outline of this algorithm, we need a few more definitions. The *interior* of a path P is the set of vertices that have degree two in P, and is denoted by P^*. If u and v are two vertices in a path P, then u-P-v denotes the subpath of P whose ends are u and v, and if u and v are consecutive then we may denote u-P-v by just u-v. For a subgraph X of G we say that $v \in V(G) \setminus V(X)$ has a neighbor in X if there exists a vertex in $V(X)$ adjacent to v.

For $A \subseteq V(G)$, a vertex is called A-*complete* if it is adjacent to every vertex in A. Two subsets $A, B \subseteq V(G)$ are *anticomplete* to each other if no vertex of A is adjacent to any vertex of B. For any two vertices a and b, let $N(a)$ be the set of neighbors of a in G, and let $N(a,b)$ be the set of all $\{a,b\}$-complete vertices in G.

Let K be a prism. Label the triangles of K $\{a_1, a_2, a_3\}$ and $\{b_1, b_2, b_3\}$ and for $i = 1,2,3$ let P_i be the path from a_i to b_i. For $i = 1,2,3$ let m_i be the vertex in P_i such that for the paths a_i-P_i-m_i, denoted by S_i, and m_i-P_i-b_i, denoted by T_i, we have $|E(T_i)| \leq |E(S_i)| \leq |E(T_i)| + 1$. A K-*major* vertex, or just a *major* vertex when there is no danger of confusion, is a vertex not in $V(K)$ whose neighbors in K are not a subset of the vertex set of a 3-vertex path contained in K.

Finally, a subgraph H is *smaller* than a subgraph H' if $|V(H)| < |V(H')|$. An induced subgraph X of G that is either a theta or a prism is called *smallest* if no other theta or prism in G is smaller than X.

The idea of the algorithm is as follows. Let G be the graph input to the algorithm and suppose there is a prism K in G. A shortest path Q with the same ends as S_1 is called a *shortcut* across K. If the subgraph induced by $V(K) \cup Q^* \setminus S_1^*$ is also a prism, then Q is *good*, and if not then Q is *bad*. Shortest paths with the same ends as one of T_1, S_2, T_2, S_3 and T_3 are also called shortcuts across K. Shortcuts will be defined in the same way for thetas. A theta or prism is *clean* if all shortcuts across it are good. The algorithm uses a procedure called *cleaning*, first introduced by Conforti and Rao in [3]. The idea of cleaning is to identify a *cleaner*, which is a subset $X \subseteq V(G)$ such that if G contains a prism or a theta, then the graph G' induced by G on $V(G) \setminus X$ contains a clean prism or a clean theta.

In G', it becomes easy to detect the clean theta or prism. For example, if G' contains a clean prism K, then for any shortcut Q across the path S_1 of K, the subgraph induced on $V(K) \cup Q^* \setminus S_1^*$ is also a prism, because Q is a good shortcut. The same is true for any other shortcut across K. So the algorithm checks all 9-tuples of vertices $a_1, a_2, a_3, m_1, m_2, m_3, b_1, b_2, b_3$; for each of the pairs a_1m_1, a_2m_2, a_3m_3, m_1b_1, m_2b_2, m_3b_3, it finds a shortest path joining the pair, then tests whether the subgraph induced by the union of these paths is a prism. If there is a clean smallest prism in G', then it can be shown that a prism is found when considering one of the 9-tuples. A similar procedure works to find clean smallest thetas.

There are three stages of cleaning required. The first two cleaning algorithms output a list of polynomially many subsets of $V(G)$, one of which is a cleaner if G contains a smallest prism or theta. The first outputs a cleaner when G contains a smallest theta, and is given in Section 2. The second outputs a cleaner when G contains a smallest prism K, and is given in Section 3. However, the second algorithm requires that the input graph G contain no K-major vertices. For this reason, we need the third cleaning algorithm, which outputs a list of subsets of which one contains all K-major vertices and is disjoint from $V(K)$. This algorithm is given in Section 4. Finally, the algorithms for detecting clean prisms and thetas are described in Section 5.

2 Theta shortcuts

Let K be a theta in G with ends a and b and paths P_1, P_2 and P_3 between them. For $1 \le i \le 3$, let s_i be the neighbor of a in P_i and t_i the neighbor of b in P_i. Also, for $1 \le i \le 3$ define $m_i \in P_i^*$ and subpaths S_i and T_i of P_i such that $S_i = s_i\text{-}P_i\text{-}m_i$, $T_i = m_i\text{-}P_i\text{-}t_i$, and $|V(T_i)| \le |V(S_i)| \le |V(T_i)| + 1$. We call P_1, P_2 and P_3 the *paths* of K and S_1, T_1, S_2, T_2, S_3, and T_3 the *half-paths* of K. The *frame* of the theta K is the 11-tuple

$$\{a, b, s_1, s_2, s_3, m_1, m_2, m_3, t_1, t_2, t_3\},$$

and we say that K has a *tidy frame* if

- a and b have no neighbors in $V(G) \setminus V(K)$, and
- for $i = 1, 2, 3$, if m_i is adjacent to a and b then m_i has no neighbors in $V(G) \setminus V(K)$.

If S is one of the half-paths of K, then a *shortcut* across S is a shortest path in G with the same ends as S. A shortcut across K is a shortcut across any half-path of K. A shortcut S' across S is *good* if no interior vertex of S' has neighbors in $V(K) \setminus V(S)$, and *bad* otherwise.

2.1. *Let K be a smallest theta in G. Assume K has a tidy frame. Then there is a vertex $c \in P_2^* \cup P_3^*$ such that c has a neighbor in the interior of every bad shortcut across S_1 or T_1.*

Proof. Let C be the cycle induced by G on $V(P_2) \cup V(P_3)$. Let Q be a bad shortcut across S_1. Suppose there is no edge between $V(Q)$ and $V(C)$ except s_1-a. Since Q is a bad shortcut, it follows that some vertex of Q has a neighbor in $V(T_1) \setminus \{m_1\}$ and then the graph induced by G on $V(C) \cup V(Q) \cup V(T_1)$ contains a theta smaller than K, a contradiction. Since m_1 has no neighbors in C and a is the only neighbor of s_1 in C, this proves that some vertex in Q^* has a neighbor in C. Let $q \in Q^*$ be the vertex with neighbors in C such that no other vertex in q-Q-m_1 has neighbors in C.

(1) *The vertex q has exactly two neighbors in C, and they are adjacent to each other.*

If $m_1 = t_1$, then any shortcut between m_1 and s_1 has empty interior, and therefore is a good shortcut. Consequently, we may assume $m_1 \neq t_1$.

Suppose q has two neighbors in C that are not adjacent to each other. Let u and v be the neighbors of q in C such that a path between u and v in C contains b and no other neighbor of q. Then u and v are not adjacent to each other. Let C_u be the path of C from u to b such that $v \notin V(C_u)$ and let C_v be the path of C from v to b such that $u \notin V(C_v)$. The paths q-u-C_u-b, q-v-C_v-b and a path between q and b with interior in $V(q\text{-}Q\text{-}m_1\text{-}T_1\text{-}b)$ form a theta, K'. Since $V(K') \subseteq (V(K) \setminus V(S_1)) \cup (V(Q) \setminus \{s_1\})$, and $|V(Q)| \leq |V(S_1)|$, we have $|V(K')| < |V(K)|$, contradicting the fact that no theta in G is smaller than K. Therefore all neighbors of q in C are pairwise adjacent.

Suppose q has exactly one neighbor in C; call it u. If u is not adjacent to b then the two paths between u and b in C and a path with interior in $V(u\text{-}q\text{-}Q\text{-}m_1\text{-}T_1\text{-}b)$ form a theta, K'. Since $V(K') \subseteq (V(K) \setminus V(S_1)) \cup (V(Q) \setminus \{s_1\})$, K' is smaller than K, a contradiction. Therefore u is adjacent to b. Assume u is not adjacent to a. Since $m_1 \neq t_1$, the two paths between u and a in C and a path P' between u and a with interior in $V(u\text{-}q\text{-}Q\text{-}m_1\text{-}S_1\text{-}s_1\text{-}a)$ form a theta, K'. If q is non-adjacent to s_1, then $|V(q\text{-}Q\text{-}m_1)| < |V(T_1)|$, and therefore $|V(K')| < |V(K)|$; and if q is adajcent to s_1, then $V(P') = \{q, s_1, a\}$, and so $V(K') = V(C) \cup \{s_1\}$ and, again, $|V(K')| < |V(K)|$, in both cases a contradiction. Therefore u is adjacent to a. Since u is adjacent to both a and b, it is one of m_2 or m_3, contradicting the fact that K has a tidy frame. Therefore, q has two neighbors in C, and they are adjacent. This proves (1).

Let R be a bad shortcut across T_1. Suppose there is no edge between $V(R)$ and $V(C)$ except t_1-b. Since R is a bad shortcut, it follows that some vertex of R has a neighbor in $V(S_1) \setminus \{m_1\}$ and then the graph induced by G on $V(C) \cup V(R) \cup V(S_1)$ contains a theta smaller than K, a contradiction. Since m_1 has no neighbors in C and b is the only neighbor of t_1 in C, this proves that some vertex in R^* has a neighbor in C. Let $r \in R^*$ be the vertex with neighbors in C such that no other vertex in r-R-m_1 has neighbors in C.

(2) *The vertex r has exactly two neighbors in C, and they are adjacent to each other.*

This proof is very similar to that of (1). If $m_1 = s_1$, then $m_1 = t_1$ and there are no shortcuts between m_1 and t_1. Therefore we may assume $m_1 \neq s_1$.

Suppose r has two neighbors in C that are not adjacent to each other. Let u and v be the neighbors of r in C such that a path between u and v in C contains a and no other neighbor of r. Then u and v are not adjacent to each other. Let C_u be the path of C from u to a such that $v \notin V(C_u)$ and let C_v be the path of C from v to a such that $u \notin V(C_v)$. The paths r-u-C_u-a, r-v-C_v-a and a path between r and a with interior contained in $V(r\text{-}R\text{-}m_1\text{-}S_1\text{-}a)$ form a theta, K'. Since $V(K') \subseteq (V(K) \setminus V(T_1)) \cup (V(R) \setminus \{t_1\})$, and $|V(R)| \leq |V(T_1)|$, we have $|V(K')| < |V(K)|$, contradicting the fact that no theta in G is smaller than K. Therefore all neighbors of r in C are pairwise adjacent.

Suppose r has exactly one neighbor in C; call it u. If u is not adjacent to a then the two paths between u and a in C and a path with interior in $V(u\text{-}r\text{-}R\text{-}m_1\text{-}S_1\text{-}a)$ form a theta, K'. Since $V(K') \subseteq (V(K) \setminus V(T_1)) \cup (V(R) \setminus \{t_1\})$, K' is smaller than K, a contradiction. Therefore u is adjacent to a. Assume u is not adjacent to b. Since $m_1 \neq s_1$, the two paths between u and b in C and a path with interior in $V(u\text{-}r\text{-}R\text{-}m_1\text{-}T_1\text{-}t_1\text{-}b)$ form a theta, K'. Since r-R-m_1 contains fewer vertices than R and $|V(R)| \leq |V(S_1)|$, we have $|V(K')| < |V(K)|$, a contradiction. Therefore u is adjacent to b. Since u is adjacent to both a and b, it is one of m_2 or m_3, contradicting the fact that K has a tidy frame. Therefore, r has two neighbors in C, and they are adjacent. This proves (2).

Let Q' be another bad shortcut across S_1, and let $q' \in Q'^*$ be the vertex with neighbors in C such that no other vertex in q'-Q'-m_1 has neighbors in C. Denote by x_q and y_q the neighbors of q in C, and by x'_q and y'_q the neighbors of q' in C.

Let R' be another bad shortcut across T_1, and let $r' \in R'^*$ be the vertex with neighbors in C such that no other vertex in r'-R'-m_1 has neighbors in C. Denote by x_r and y_r the neighbors of r in C, and by x'_r and y'_r the neighbors of r' in C.

(3) *No two of the subsets $\{x_q, y_q\}$, $\{x'_q, y'_q\}$, $\{x_r, y_r\}$ and $\{x'_r, y'_r\}$ are disjoint.*

Assume that $\{x_q, y_q\}$ is disjoint from $\{x'_q, y'_q\}$. Since K has a tidy frame, $\{a, b\} \cap \{x_q, x'_q, y_q, y'_q\} = \emptyset$. Let Q'' be the path between q and q' with interior in $V(q\text{-}Q\text{-}m_1) \cup V(q'\text{-}Q'\text{-}m_1)$. The triangles $\{q, x_q, y_q\}$ and $\{q', x'_q, y'_q\}$, the two paths between them in C and the path Q'' form a prism, K'. Since

$$|V(Q'')| \leq 2|V(S_1)| - 3 < |V(S_1)| + |V(T_1)| - 1,$$

and $V(K') \subseteq V(C) \cup V(Q'')$, we deduce that K' is smaller than K, a contradiction. Therefore, $\{x_q, y_q\} \cap \{x'_q, y'_q\} \neq \emptyset$.

Assume that $\{x_r, y_r\}$ is disjoint from $\{x'_r, y'_r\}$. Since K has a tidy frame, $\{a, b\} \cap \{x_r, x'_r, y_r, y'_r\} = \emptyset$. Let R'' be the path between r and r' with interior in $V(r\text{-}R\text{-}m_1) \cup V(r'\text{-}R'\text{-}m_1)$. The triangles $\{r, x_r, y_r\}$ and $\{r', x'_r, y'_r\}$, the two paths between them in C and the path R'' form a prism, K'. Since

$$|V(R'')| \leq 2|V(T_1)| - 3 < |V(S_1)| + |V(T_1)| - 2,$$

and $V(K') \subseteq V(C) \cup V(R'')$, we deduce that K' is smaller than K, a contradiction. Therefore, $\{x_r, y_r\} \cap \{x'_r, y'_r\} \neq \emptyset$.

Suppose $\{x_q, y_q\}$ is disjoint from $\{x_r, y_r\}$. Then the triangles $\{q, x_q, y_q\}$ and $\{r, x_r, y_r\}$, the two paths in C between $\{x_q, y_q\}$ and $\{x_r, y_r\}$ and the path between q and r with interior in $V(q\text{-}Q\text{-}m_1\text{-}R\text{-}r)$ form a prism, K'. Since $V(K') \subseteq V(C) \cup V(Q) \cup V(R) \setminus \{s_1, t_1\}$, the prism K' is smaller than K, a contradiction. Therefore $\{x_q, y_q\} \cap \{x_r, y_r\} \neq \emptyset$. By symmetry, we also have that $\{x_q, y_q\} \cap \{x'_r, y'_r\} \neq \emptyset$, $\{x'_q, y'_q\} \cap \{x_r, y_r\} \neq \emptyset$, and $\{x'_q, y'_q\} \cap \{x'_r, y'_r\} \neq \emptyset$. This proves (3).

Let U, W and Z each be a bad shortcut across either S_1 or T_1. Since the choices of Q, Q', R and R' were arbitrary, it follows from (1), (2) and (3) that there exist $u \in U^*$, $w \in W^*$ and $z \in Z^*$ such that u has two adjacent neighbors $x_u, y_u \in V(C)$, w has two adjacent neighbors $x_w, y_w \in V(C)$, z has two adjacent neighbors $x_z, y_z \in V(C)$, and no two of $\{x_u, y_u\}$, $\{x_w, y_w\}$ and $\{x_z, y_z\}$ are disjoint.

Assume that $\{x_w, y_w\} \cap \{x_z, y_z\} = \{x_w\}$ and $\{x_w, y_w\} \cap \{x_u, y_u\} = \{y_w\}$. Then, since x_w and y_w are adjacent and x_z and y_z are adjacent, $\{x_z, y_z\}$

contains x_w and the vertex adjacent to x_w in $V(C) \setminus \{y_w\}$. Similarly, since x_u and y_u are adjacent, $\{x_u, y_u\}$ contains y_w and the vertex adjacent to y_w in $V(C) \setminus \{x_w\}$. Since x_w and y_w have no common neighbors in C, we have that $\{x_z, y_z\} \cap \{x_u, y_u\} = \emptyset$, a contradiction. Therefore, by symmetry between x_w and y_w we may assume that $x_w \in \{x_z, y_z\}$ and $x_w \in \{x_u, y_u\}$. Therefore every bad shortcut across S_1 or T_1 contains an interior vertex adjacent to x_w; taking $c = x_w$ proves 2.1. ☐

2.2. *There is an algorithm with the following specifications:*

- **Input:** *A graph G.*
- **Output:** *A sequence of subsets $X_1, \dots, X_r$ of $V(G)$ with $r \leq |V(G)|^3$, such that for every smallest theta K in G with a tidy frame and $i \in \{1, 2, 3\}$, one of $X_1, \dots, X_r$ is disjoint from $V(K)$ and contains a vertex of every bad shortcut across the half-paths S_i and T_i of K.*
- **Running Time:** *$O(|V(G)|^4)$.*

Proof. The algorithm is as follows. Enumerate all triples of vertices (c, u, v) for which u and v are adjacent to c. For each such triple, compute the subset $N(c) \setminus \{u, v\}$. Let $X_1, \dots, X_r$ be the subsets generated. We output the list $\emptyset, X_1, \dots, X_r$. That concludes the description of the algorithm; we now prove that it works correctly. The number of subsets generated is at most $|V(G)|^3$ and the running time is $O(|V(G)|^4)$ because it takes linear time to compute $N(c)$. It remains to check that the sequence $X_1, \dots, X_r$ has the properties claimed. Let K be a smallest theta in G and assume it has a tidy frame. By 2.1, for every $i = 1, 2, 3$ there is a vertex $c' \in V(K)$ such that every bad shortcut across S_i or T_i contains an interior vertex adjacent to c'. The vertex c' has exactly two neighbors in K; call them u' and v'. One of the triples enumerated by the algorithm will be (c', u', v'). The subset $N(c') \setminus \{u', v'\}$ includes all neighbors of c' except its two neighbors that are in K, and is therefore disjoint from $V(K)$. This proves 2.2. ☐

3 Prism shortcuts

Let K be a prism in G. Label the triangles of K $\{a_1, a_2, a_3\}$ and $\{b_1, b_2, b_3\}$ and the paths P_1, P_2 and P_3 such that for $i = 1, 2, 3$ the path P_i has ends a_i and b_i. For $i = 1, 2, 3$ choose $m_i \in P_i^*$ and subpaths S_i and T_i such that $S_i = a_i\text{-}P_i\text{-}m_i$, $T_i = m_i\text{-}P_i\text{-}b_i$ and $|V(S_i)| \leq |V(T_i)| + 1$ and $|V(T_i)| \leq |V(S_i)| + 1$. We call P_1, P_2 and P_3 the *paths* of K and S_1, T_1, S_2,

T_2, S_3, and T_3 the *half-paths* of K. The 9-tuple

$$\{a_1, a_2, a_3, m_1, m_2, m_3, b_1, b_2, b_3\}$$

is the *frame* of K. A prism K has a *tidy frame* if for $i = 1, 2, 3$,

- neither a_i nor b_i has neighbors in $V(G) \setminus V(K)$, and
- if m_i is adjacent to a_i and b_i, then m_i has no neighbors in $V(G) \setminus V(K)$.

The prism K in G is *almost clean* if it has a tidy frame and G contains no K-major vertices.

If S is one of the half-paths of K, then a *shortcut* across S is a shortest path in G with the same ends as S. A shortcut across K is a shortcut across some half-path of K. A shortcut Q across S is *good* if no interior vertex of Q has neighbors in $V(K) \setminus V(S)$, and *bad* otherwise.

Let C be the cycle induced by G on $V(P_2) \cup V(P_3)$.

3.1. *Let K be a smallest prism. For every bad shortcut Q across S_1, some vertex of Q^* has a neighbor in C.*

Proof. Suppose no vertex of Q^* has a neighbor in C. Then, since Q is bad, some vertex of Q^* has a neighbor in $V(T_1) \setminus \{m_1\}$. Therefore, there exists a path P'_1 from a_1 to b_1 with $P'^*_1 \subset V(Q) \cup V(T_1) \setminus \{m_1\}$. So there is a prism contained in $V(C) \cup V(P'_1)$ that is smaller than K, a contradiction because K is a smallest prism. This proves 3.1. ☐

In view of 3.1, for a bad shortcut Q across S_1 or T_1, define $\beta_Q \in Q^*$ to be the vertex with neighbors in $V(C)$ such that no other vertex in β_Q-Q-m_1 has neighbors in $V(C)$.

3.2. *Let K be an almost clean smallest prism. For a bad shortcut Q across S_1, either $|N(\beta_Q) \cap V(K)| = |N(\beta_Q) \cap V(C)| = 1$ or $N(\beta_Q) \cap V(K)$ is the vertex set of a 3-vertex path of C.*

Proof. Since K is almost clean, β_Q is not a K-major vertex, and therefore $N(\beta_Q) \cap V(K) \subseteq N(\beta_Q) \cap V(C)$. Assume β_Q has exactly two neighbors in C, u and v. If u and v are not adjacent, then u-β_Q-v and the two paths in C between u and v form a theta that has fewer vertices than K, a contradiction. So u and v are adjacent. Denote β_Q-Q-m_1 by U. Since K is almost clean, $\{u, v\} \cap \{b_1, b_2, b_3\} = \emptyset$, so there is a prism K' containing the triangles $\{b_1, b_2, b_3\}$ and $\{\beta_Q, u, v\}$ with $V(K') \subseteq V(C) \cup V(T_1) \cup V(U)$ (the three paths of K' are two paths of C and the path contained in $V(U) \cup V(T_1)$). Since $|V(U)| < |V(S_1)|$ (because $a_1 \notin V(U)$), K' is smaller than K, a contradiction. Therefore, β_Q has either a unique neighbor in C or at least three neighbors in C. Since K is almost clean, β_Q is not K-major, so if it has at least three neighbors in C then it has exactly three neighbors in C, which are the vertex set of a 3-vertex path of C. This proves 3.2. $\square$

3.3. *Let K be a smallest prism in G. Assume K is almost clean. Then there exist two edges e_1 and e_2, one in $E(P_2)$ and the other in $E(P_3)$, such that either*

- *there exists a bad shortcut T across S_1 with $|N(\beta_T) \cap V(C)| = 1$ and for every bad shortcut Q across S_1, if $|N(\beta_Q) \cap V(C)| = 1$ then β_Q is adjacent to an end of e_1 and if $|N(\beta_Q) \cap V(C)| > 1$ then β_Q is adjacent to both ends of e_1, or*
- *for every bad shortcut Q, $|N(\beta_Q) \cap V(C)| = 3$ and β_Q is adjacent either to both ends of e_1 or to both ends of e_2.*

Proof. We start with a few observations about bad shortcuts across S_1.

(1) *There exists an edge $e_1 \in E(C)$ such that for every bad shortcut Q across S_1, if $|N(\beta_Q) \cap V(C)| = 1$ then β_Q is adjacent to an end of e_1.*

We may assume that there exist two distinct bad shortcuts Q and Q' across S_1 such that β_Q has a unique neighbor x in C, $\beta_{Q'}$ has a unique neighbor y in C, and $x \neq y$. Let U be the path β_Q-Q-m_1, let U' be the path $\beta_{Q'}$-Q'-m_1 and let W be a path between β_Q and $\beta_{Q'}$ with interior in $V(U) \cup V(U')$. If x and y are nonadjacent, then the two paths between them in C and the path x-β_Q-W-$\beta_{Q'}$-y form a theta, K'. Since $|V(W)| \leq 2|V(S_1)| - 3$ and $|V(P_1)| \geq 2|V(S_1)| - 2$, K' is smaller than K, a contradiction. Therefore, x is adjacent to y.

Let Q'' be any other bad shortcut across S_1 such that $\beta_{Q''}$ has a unique neighbor z in C. By the previous argument, applied to the pairs Q, Q'' and Q', Q'', it follows that z is equal or adjacent to x and equal or adjacent to y. Then either $z \in \{x, y\}$ or z is adjacent to both x and y. But since x and y are adjacent vertices of C, they have no common neighbor in C, so $z \in \{x, y\}$. Therefore, there exists an edge $e_1 = x$-y such that whenever $\beta_{Q''}$ has a unique neighbor in C, $\beta_{Q''}$ is adjacent to an end of e_1. This proves (1).

(2) *There exist edges $e_1 \in E(P_2)$ and $e_2 \in E(P_3)$ such that for every bad shortcut Q across S_1, if $|N(\beta_Q) \cap V(C)| = 3$ then β_Q is either adjacent to both ends of e_1 or to both ends of e_2.*

Let Q be a bad shortcut across S_1 such that $\beta_{Q'}$ has three neighbors in C. By 3.2, $N(\beta_Q) \cap V(C)$ is the vertex set of a 3-vertex path of C. Since K is almost clean, $N(\beta_Q) \cap V(C)$ is contained in either P_2^* or P_3^*.

Therefore, in order to prove (2), it is enough to show that there exist edges $e_1 \in E(P_2)$ and $e_2 \in E(P_3)$ such that for every bad shortcut Q across S_1, if $|N(\beta_Q) \cap V(P_2)| = 3$ then β_Q is adjacent to both ends of e_1, and if $|N(\beta_Q) \cap V(P_3)| = 3$, then β_Q is adjacent to both ends of e_2.

Let Q and Q' be bad shortcuts across S_1. Suppose that each of β_Q and $\beta_{Q'}$ has three neighbors in P_2. Then each of $N(\beta_Q) \cap V(C)$ and $N(\beta_{Q'}) \cap V(C)$ is the vertex set of a 3-vertex path of P_2.

We claim that $|N(\beta_Q) \cap V(C) \cap N(\beta_{Q'})| \geq 2$. Assume that $|N(\beta_Q) \cap V(C) \cap N(\beta_{Q'})| < 2$. If β_Q and $\beta_{Q'}$ are adjacent to each other, then there is a path P_2' between a_2 and b_2 whose interior is contained in $P_2^* \cup \{\beta_Q, \beta_{Q'}\}$ and is anticomplete to $V(K) \setminus P_2^*$. Then $(V(K) \setminus P_2^*) \cup P_2'^*$ contains a prism smaller than K, a contradiction. So we may assume that β_Q and $\beta_{Q'}$ are not adjacent to each other. Then there are two paths A and B between β_Q and $\beta_{Q'}$ whose interiors are anticomplete to each other and are contained in $V(C)$. The paths A, B and a path between β_Q and $\beta_{Q'}$ with interior in $V(Q) \cup V(Q')$ form a theta, K'. Since $|V(Q) \cup V(Q')| \leq 2|V(S_1)| - 3 < |V(S_1)| + |V(T_1)| - 1 = |V(P_1)|$, the theta K' is smaller than the prism K, a contradiction. This proves the claim.

Label the neighbors of β_Q in C x, y and z and the neighbors of $\beta_{Q'}$ in C x', y' and z' such that x-y-z and x'-y'-z' are paths of C. Let Q'' be another bad shortcut across S_1 such that $\beta_{Q''}$ has three neighbors x'', y'' and z'' in P_2 and x''-y''-z'' is a path of P_2. Then we have that $|\{x'', y'', z''\} \cap \{x, y, z\}| \geq 2$ and $|\{x'', y'', z''\} \cap \{x', y', z'\}| \geq 2$. Since x-y-z is a path, either $\{x, y\} \subset \{x', y', z'\}$ or $\{y, z\} \subset \{x', y', z'\}$, and similarly, either $\{x, y\} \subset \{x'', y'', z''\}$ or $\{y, z\} \subset \{x'', y'', z''\}$. Assume that $\{x, y, z\} \cap \{x', y', z'\} = \{x, y\}$ and $\{x, y, z\} \cap \{x'', y'', z''\} = \{y, z\}$. Then since x-y-z is a path and x'-y'-z' is a path, $\{x', y', z'\}$ contains $\{x, y\}$ and the vertex adjacent to x in $V(P_2) \setminus \{y\}$. Similarly, since x''-y''-z'' is a path, $\{x'', y'', z''\}$ contains $\{y, z\}$ and the vertex adjacent to z in $V(P_2) \setminus \{y\}$. Since x and z are not adjacent and x and z have no common neighbors in $V(P_2) \setminus \{y\}$, it follows that $\{x', y', z'\} \cap \{x'', y'', z''\} = \{y\}$, a contradiction. Therefore, by symmetry between $\{x, y\}$ and $\{y, z\}$ we may assume that $\{x, y\} \subset \{x', y', z'\}$ and $\{x, y\} \subset \{x'', y'', z''\}$. This proves that for every bad shortcut R across S_1, if $|N(\beta_R) \cap V(P_2)| = 3$ then β_R is adjacent to both ends of $e_1 = x$-y.

From the symmetry, there exists an edge $e_2 \in E(P_3)$ such that for every bad shortcut R across S_1, if $|N(\beta_R) \cap V(P_3)| = 3$ then β_R is adjacent to both ends of e_2. This proves (2).

(3) *If Q and Q' are bad shortcuts across S_1 with $|N(\beta_Q) \cap V(C)| = 1$ and $|N(\beta_{Q'}) \cap V(C)| = 3$, then the neighbor of β_Q in C is adjacent to $\beta_{Q'}$.*

Denote by x the neighbor of β_Q in C. By 3.2, the neighbors of $\beta_{Q'}$ in C are the vertex set of a 3-vertex path of C, say c_1-c_2-c_3.

Assume that $x \notin \{c_1, c_2, c_3\}$. Then the subgraph induced by G on $(V(C) \setminus \{c_2\}) \cup V(Q) \cup V(Q')$ contains a theta, K' (the three paths of K' are two paths between $\beta_{Q'}$ and x with interior in C and the path with interior in $V(\beta_{Q'}$-Q'-m_1-Q-β_Q-$x)$). Since $V(K') \subseteq V(C) \cup Q^* \cup Q'^* \cup \{m_1\}$ and $|Q^*| + |Q'^*| + 1 \leq 2|V(S_1)| - 3 < |V(S_1)| + |V(T_1)| - 1$, the theta K' is smaller than K, a contradiction. Therefore, $x \in \{c_1, c_2, c_3\}$. This proves (3).

If for every bad shortcut Q, $|N(\beta_Q) \cap V(C)| = 3$, then the second outcome of the theorem holds by (2). So we may assume that there exists a bad shortcut Q with $N(\beta_Q) \cap V(C) = \{x\}$. Suppose there exists a bad shortcut Q' with $|N(\beta_{Q'}) \cap V(C)| = \{y\}$ such that $x \neq y$. Then by (1) and (3) the first outcome of the theorem holds with e_1 being the edge as in (1). So we may assume that for every bad shortcut Q' with $|N(\beta_{Q'}) \cap V(C)| = 1$, $\beta_{Q'}$ is adjacent to x. Let c_1-x-c_2 be the 3-vertex path in C with x as its interior vertex. If there exist bad shortcuts Q' and Q'' with $|N(\beta_{Q'}) \cap V(C)| = 3$ and $|N(\beta_{Q''}) \cap V(C)| = 3$ such that $\beta_{Q'}$ is not adjacent to c_1 and $\beta_{Q''}$ is not adjacent to c_2, then $N(\beta_{Q'}) \cap V(C) \cap N(\beta_{Q''}) = \{x\}$, contradicting (2). Thus one of c_1 and c_2, say c_1, is complete to $\{\beta_{Q'}, \beta_{Q''}\}$, and the theorem holds with $e_1 = c_1$-x. This completes the proof of 3.3. □

3.4. *Let K be a smallest prism in G. Assume K is almost clean. Then there exist two edges e_1 and e_2, one in $E(P_2)$ and the other in $E(P_3)$, such that either*

- *there exists a bad shortcut T across S_1 or T_1 with $|N(\beta_T) \cap V(C)| = 1$, and for every bad shortcut Q across S_1 or T_1, if $|N(\beta_Q) \cap V(C)| = 1$ then β_Q is adjacent to an end of e_1 and if $|N(\beta_Q) \cap V(C)| > 1$ then β_Q is adjacent to both ends of e_1, or*
- *for every bad shortcut Q across S_1 or T_1, $|N(\beta_Q) \cap V(C)| = 3$ and β_Q is adjacent either to both ends of e_1 or to both ends of e_2.*

Proof. We start with a few observations about bad shortcuts across S_1 and T_1.

(1) *Let Q be a bad shortcut across S_1 with $|N(\beta_Q) \cap V(C)| = 1$ and R a bad shortcut across T_1 with $|N(\beta_R) \cap V(C)| = 1$. Then the neighbor of β_Q in C and the neighbor of β_R in C are either equal or adjacent to each other.*

Let $\{x_q\} = N(\beta_Q) \cap V(C)$ and $\{x_r\} = N(\beta_R) \cap V(C)$. Assume that x_q and x_r are distinct and nonadjacent. There are two paths between them, A and B, each of length at least two and contained in C. There is a path P between x_q and x_r with interior in $V(\beta_Q\text{-}Q\text{-}m_1\text{-}R\text{-}\beta_R)$. The three paths A, B and P form a theta, K'. Since $|V(P)| \leq |Q^* \cup \{m_1\} \cup R^*| < |V(P_1)|$, the theta K' is smaller than the prism K, a contradiction. This proves (1).

(2) *Let Q be a bad shortcut across S_1 with $|N(\beta_Q) \cap V(C)| = 3$ and $N(\beta_Q) \cap V(C) \subseteq P_2^*$. Let R be a bad shortcut across T_1 with $|N(\beta_R) \cap V(C)| = 3$ and $N(\beta_R) \cap V(C) \subseteq P_2^*$. Then there is an edge in $E(P_2)$ whose ends are $\{\beta_Q, \beta_R\}$-complete.*

Assume that $|N(\beta_Q) \cap V(C) \cap N(\beta_R)| < 2$. By 3.2, each of $N(\beta_Q) \cap V(K)$ and $N(\beta_R) \cap V(K)$ is the vertex set of a 3-vertex path of C. If β_Q is adjacent to β_R then there is a path A between a_2 and b_2 with interior in $P_2^* \cup \{\beta_Q, \beta_R\}$ such that $\{\beta_Q, \beta_R\} \subset V(A)$. Since $|N(\beta_Q) \cap V(C)| = |N(\beta_R) \cap V(C)| = 3$, it follows that $|(N(\beta_Q) \cup N(\beta_R)) \cap V(C)| \geq 5$. So $|V(P_2) \setminus V(A)| \geq 3$ and $|V(A) \setminus V(P_2)| = 2$, which implies that $|V(A)| < |V(P_2)|$. It follows from 3.2 that A^* is anticomplete to P_1^* and P_3^*, and so the induced subgraph with vertex set $(V(K) \setminus V(P_2)) \cup V(A)$ is a prism smaller than K, a contradiction. This proves that β_Q is not adjacent to β_R. It follows that there are two paths B and B' between β_Q and β_R whose interiors are anticomplete to each other and are contained in $V(C)$. There is also a path P between them with interior contained in $V(\beta_Q\text{-}Q\text{-}m_1\text{-}R\text{-}\beta_R)$. Since β_Q is not adjacent to β_R, P has length at least two. Since $|V(P)| \leq |Q^* \cup \{m_1\} \cup R^*| < |V(P_1)|$, the paths B, B' and P form a theta smaller than K, a contradiction. This proves that $|N(\beta_Q) \cap V(C) \cap N(\beta_R)| \geq 2$. Since each of $N(\beta_Q) \cap V(C)$ and $N(\beta_R) \cap V(C)$ is the vertex set of a 3-vertex subpath of P_2, (2) follows.

(3) *Let Q be a bad shortcut across S_1 with $|N(\beta_Q) \cap V(C)| = 1$ and R a bad shortcut across T_1 with $|N(\beta_R) \cap V(C)| = 3$. Then the neighbor of β_Q in C is adjacent to β_R.*

Let $\{x_q\} = N(\beta_Q) \cap V(C)$. Assume that $x_q \notin N(\beta_R)$. There are two paths A and B between x_q and β_R whose interiors are anticomplete to each other and are contained in $V(C)$. There is a path P between x_q and β_R with interior in $V(\beta_Q\text{-}Q\text{-}m_1\text{-}R\text{-}\beta_R)$. Since $|V(P)| \leq |Q^* \cup \{m_1\} \cup R^*| + 1 < |V(P_1)|$, the paths A, B and P form a theta smaller than K, a contradiction. This proves (3).

(4) *Suppose that there exist* $p \in V(P_2)$ *and* $s \in V(P_3)$ *such that for every bad shortcut* Q *across* S_1 *or* T_1, β_Q *is adjacent to* p *or* s. *Then the theorem holds.*

Since there are no K-major vertices in G, for every bad shortcut Q across S_1 or T_1 with $N(\beta_Q) \cap V(C) \subseteq P_2^*$, $N(\beta_Q) \cap V(C)$ is contained in the vertex set of the 5-vertex path c_1-c_2-p-c_3-c_4 contained in C. Suppose there exist two bad shortcuts P and P' across either S_1 or T_1 with $N(\beta_P) \cap V(C) = \{c_1, c_2, p\}$ and $N(\beta_{P'}) \cap V(C) = \{p, c_3, c_4\}$. Then $N(\beta_P) \cap V(C) \cap N(\beta_{P'}) = \{p\}$. If both P and P' are shortcuts across S_1 or both are shortcuts across T_1, then this contradicts both outcomes of 3.3. If one of P and P' is a shortcut across S_1 and the other is a shortcut across T_1, then this contradicts (2). Therefore, there exists a 4-vertex path A of C with $p \in A^*$ such that for every bad shortcut P across S_1 or T_1, if $N(\beta_P) \cap V(C) \subseteq P_2^*$ then $N(\beta_P) \cap V(C) \subset V(A)$. Let e_1 be the edge with both ends in A^*. It follows that for every bad shortcut P across S_1 or T_1 with $N(\beta_P) \cap V(C) \subseteq P_2^*$, if $|N(\beta_P) \cap V(C)| = 1$ then β_P is adjacent to the end p of e_1, and if $|N(\beta_P) \cap V(C)| > 1$ then β_P is adjacent to both ends of e_1. By symmetry between P_2 and P_3, it also follows that there is an edge $e_2 \in E(P_3)$ such that for every bad shortcut P across S_1 or T_1 with $N(\beta_P) \cap V(C) \subseteq P_3^*$, if $|N(\beta_P) \cap V(C)| = 1$ then β_P is adjacent to an end of e_2, and if $|N(\beta_P) \cap V(C)| > 1$ then β_P is adjacent to both ends of e_2. So if for all bad shortcuts Q across S_1 and T_1, $|N(\beta_Q) \cap V(C)| > 1$, then the second outcome of the theorem holds. If not, then by (1) and (3), either for all bad shortcuts Q across S_1 and T_1, $N(\beta_Q) \cap V(C) \subseteq P_2^*$, or for all bad shortcuts Q across S_1 and T_1, $N(\beta_Q) \cap V(C) \subseteq P_3^*$. In both cases, (with e_2 in place of e_1 in the second case), the first outcome of the theorem holds. This proves (4).

First assume that there exist a bad shortcut Q across S_1 with $|N(\beta_Q) \cap V(C)| = 1$ and a bad shortcut R across T_1 with $|N(\beta_R) \cap V(C)| = 1$. The first outcome of 3.3 holds for both S_1 and T_1. Let $\{x_q\} = N(\beta_Q) \cap V(C)$ and $\{x_r\} = N(\beta_R) \cap V(C)$. Assume that $x_q \neq x_r$. By (1), x_q is adjacent to x_r. Let Q' be any other bad shortcut across S_1 and let R' be any other bad shortcut across T_1. It follows from (1) that if $|N(\beta_{Q'}) \cap V(C)| = 1$ then $\beta_{Q'}$ is adjacent to x_q or x_r; it also follows that if $|N(\beta_{R'}) \cap V(C)| = 1$ then $\beta_{R'}$ is adjacent to x_r or x_q. Also, 3.3 implies that if $|N(\beta_{Q'}) \cap V(C)| > 1$ then $\beta_{Q'}$ is adjacent to x_q, and if $|N(\beta_{R'}) \cap V(C)| > 1$ then $\beta_{R'}$ is adjacent to x_r. By (3), if $|N(\beta_{Q'}) \cap V(C)| > 1$ then $\beta_{Q'}$ is adjacent to x_r and if $|N(\beta_{R'}) \cap V(C)| > 1$ then $\beta_{R'}$ is adjacent to x_q. Therefore, the first outcome of the theorem holds with e_1 being x_q-x_r. Now we may assume that for every bad shortcut Q across S_1 with $|N(\beta_Q) \cap V(C)| = 1$ and R across T_1 with $|N(\beta_R) \cap V(C)| = 1$, if $\{x_q\} = N(\beta_Q) \cap V(C)$ and $\{x_r\} = N(\beta_R) \cap V(C)$, then $x_q = x_r$. By the first outcome of 3.3, for any bad shortcut Q' across S_1 with $|N(\beta_{Q'}) \cap V(C)| > 1$, $\beta_{Q'}$ is adjacent to x_q, and for any bad shortcut R' across T_1 with $|N(\beta_{R'}) \cap V(C)| > 1$, $\beta_{R'}$ is adjacent to x_r, which equals x_q. By (4) with $p = x_q$, the first outcome of the theorem holds.

Next assume that there exists a bad shortcut Q across S_1 with $|N(\beta_Q) \cap V(C)| = 1$. Then we may assume because of the previous paragraph that for every bad shortcut R across T_1, $|N(\beta_R) \cap V(C)| > 1$. The first outcome of 3.3 holds for shortcuts across S_1 and the second outcome of 3.3 holds for shortcuts across T_1. Let $\{x_q\} = N(\beta_Q) \cap V(C)$. Suppose there is a bad shortcut Q' across S_1 with $N(\beta_{Q'}) \cap V(C) = \{x_{q'}\}$ such that $x_{q'} \neq x_q$. By the first outcome of 3.3, x_q is adjacent to x'_q and for every bad shortcut Q'' across S_1, if $|N(\beta_{Q''}) \cap V(C)| = 1$ then $\beta_{Q''}$ is adjacent to either x_q or to $x_{q'}$, and if $|N(\beta_{Q''}) \cap V(C)| > 1$ then $\beta_{Q''}$ is adjacent to both x_q and $x_{q'}$. By (3), for every bad shortcut R' across T_1, both x_q and $x_{q'}$ are adjacent to $\beta_{R'}$. Then the first outcome of the theorem holds with e_1 being the edge x_q-$x_{q'}$. So for every bad shortcut Q' across S_1 with $|N(\beta_{Q'}) \cap V(C)| = 1$ we may assume that $\beta_{Q'}$ is adjacent to x_q. By the first outcome of 3.3, for any bad shortcut Q' across S_1 with $|N(\beta_{Q'}) \cap V(C)| > 1$, $\beta_{Q'}$ is adjacent to x_q. By (3), for any bad shortcut R' across T_1, $\beta_{R'}$ is adjacent to x_q. Therefore, by (4) with $p = x_q$, the first outcome of the theorem holds.

If there exists a bad shortcut R across T_1 with $|N(\beta_R) \cap V(C)| = 1$, then by symmetry, the argument of the previous paragraph applies. So we may finally assume that for every shortcut P across either S_1 or T_1, $|N(\beta_P) \cap V(C)| > 1$. Then the second outcome of 3.3 holds for both S_1 and T_1. Let Q be a bad shortcut across S_1 with $N(Q) \cap V(C) \subseteq P_2^*$. Let c_1-p-c_2 be a subpath of P_2 such that $N(Q) \cap V(C) = \{c_1, p, c_2\}$. By the second outcome of 3.3, for any bad shortcut Q' across S_1 with $N(\beta_{Q'}) \cap V(C) \subseteq P_2^*$, $\beta_{Q'}$ is adjacent to p. Also, by (2), for any bad shortcut R across T_1 with $N(\beta_R) \cap V(C) \subseteq P_2^*$, β_R is adjacent to p. By symmetry between P_2 and P_3, there is also a vertex $s \in P_3^*$ such that for every bad shortcut R across S_1 or T_1 with $N(\beta_R) \cap V(C) \subseteq P_3^*$, β_R is adjacent to s. Therefore, by (4), the second case of the theorem holds. This proves 3.4. □

3.5. *There is an algorithm with the following specifications:*

- **Input:** *A graph G.*
- **Output:** *A sequence of subsets $X_1, \ldots, X_r$ of $V(G)$ with $r \leq 2|V(G)|^4$, such that for every almost clean smallest prism K in G and $i \in$*

$\{1,2,3\}$, one of $X_1, \dots, X_r$ is disjoint from $V(K)$ and contains a vertex of every bad shortcut across the half-paths S_i and T_i of K.

- **Running Time:** $O(|V(G)|^5)$.

Proof. The algorithm is as follows. Enumerate all quadruples of distinct vertices (u_1, u_2, u_3, u_4). For each, define the subset

$$Y = (N(u_2) \cup N(u_3)) \setminus \{u_1, u_2, u_3, u_4\}$$

and the subset

$$Z = N(u_1, u_2) \cup N(u_3, u_4).$$

Let $Y_1, \dots, Y_s$ and $Z_1, \dots, Z_s$ be the subsets generated. Output $\emptyset$, $Y_1, \dots, Y_s$, $Z_1, \dots, Z_s$.

That concludes the description of the algorithm; we now prove that it works correctly. It takes time $O(|V(G)|^5)$ to find the sets $Y_1, \dots, Y_s$ and $Z_1, \dots, Z_s$ and the number of subsets in the output is $\leq 2|V(G)|^4$. It remains to check that the output sequence has the properties claimed. Let K be an almost clean smallest prism in G with half-paths S_1, T_1, S_2, T_2, S_3 and T_3. We may assume by symmetry that $i = 1$. Apply 3.4 to S_1 and T_1.

If the first outcome of 3.4 holds, then there is an edge in $P_2^* \cup P_3^*$, say u_2-u_3, such that every bad shortcut across S_1 or T_1 contains an interior vertex adjacent to either u_2 or u_3. Choose u_1 to be the neighbor of u_2 in $V(K) \setminus \{u_3\}$ and u_4 to be the neighbor of u_3 in $V(K) \setminus \{u_2\}$. Then the subset $(N(u_2) \cup N(u_3)) \setminus \{u_1, u_2, u_3, u_4\}$ is disjoint from $V(K)$ and contains an interior vertex of every bad shortcut across S_1 or T_1.

If the second outcome of 3.4 holds, then there are two edges in $P_2^* \cup P_3^*$, say u_1-u_2 and u_3-u_4, such that the set $N(u_1, u_2) \cup N(u_3, u_4)$ contains an interior vertex of every bad shortcut across S_1 or T_1. Since u_1 and u_2 are adjacent, they have no common neighbors in $V(K)$, and since u_3 and u_4 are adjacent, they have no common neighbors in $V(K)$, so $N(u_1, u_2) \cup N(u_3, u_4)$ is disjoint from $V(K)$. This proves 3.5. ☐

4 Major vertices

In this section, let K be a smallest prism in a graph G, and assume that K has a tidy frame. Let the vertices of P_1 be $f_1, f_2, \dots, f_n$, numbered in order with $f_1 = a_1$ and $f_n = b_1$. Let the vertices of P_2 be $h_1, h_2, \dots, h_m$, numbered in order with $h_1 = a_2$ and $h_m = b_2$, and let the vertices of P_3 be $g_1, g_2, \dots, g_p$, numbered in order with $g_1 = a_3$ and $g_p = b_3$.

Recall that a K-major vertex is a vertex in $V(G)\setminus V(K)$ whose neighbors in K are not contained in a 3-vertex path in K. Since K has a tidy frame, every K-major vertex is anticomplete to $\{a_1, a_2, a_3, b_1, b_2, b_3\}$. For a K-major vertex x, if x has neighbors in P_1, let j_x be minimal such that x is adjacent to f_{j_x} and let k_x be maximal such that x is adjacent to f_{k_x}. Similarly, let s_x be minimal and t_x maximal such that x is adjacent to h_{s_x} and h_{t_x}, and let c_x be minimal and d_x maximal such that x is adjacent to g_{c_x} and g_{d_x}.

We start with a few easy but useful lemmas.

4.1. *A K-major vertex has at least three neighbors in K.*

Proof. Let x be a major vertex with $N(x) \cap V(K) = \{u, v\}$. We may assume by symmetry that $\{u, v\} \subset V(P_2) \cup V(P_3)$. Since x is major, u and v are not adjacent, so the subgraph induced by G on $V(P_2) \cup V(P_3) \cup \{x\}$ is a theta smaller than K, a contradiction. □

4.2. *There is no K-major vertex with exactly two neighbors in P_2 and no neighbors in P_3.*

Proof. Let x be a K-major vertex with $N(x) \cap V(P_2) = \{h_{s_x}, h_{t_x}\}$ and $N(x) \cap V(P_3) = \emptyset$. By 4.1, x has at least one neighbor in P_1. If $s_x + 1 < t_x$, then the subgraph induced by G on $V(P_2) \cup V(P_3) \cup \{x\}$ is a theta smaller than K, a contradiction. So we may assume that $s_x + 1 = t_x$. Now the triangles $\{h_{s_x}, h_{t_x}, x\}$ and $\{a_1, a_2, a_3\}$ and the paths h_{s_x}-P_2-a_2, h_{t_x}-P_2-b_2-b_3-P_3-a_3 and x-f_{j_x}-P_1-a_1 form a prism, K'. Since $V(K') \subseteq V(K) \cup \{x\} \setminus \{f_{j_x+1}, b_1\}$, the prism K' is smaller than K unless $f_{j_x+1} = b_1$. Also, the triangles $\{x, h_{s_x}, h_{t_x}\}$ and $\{b_1, b_2, b_3\}$ and the paths h_{s_x}-P_2-a_2-a_3-P_3-b_3, h_{t_x}-P_2-b_2 and x-f_{k_x}-P_1-b_1 form a prism, K''. Since $V(K'') \subseteq V(K) \cup \{x\} \setminus \{f_{k_x-1}, a_1\}$, the prism K'' is smaller than K unless $f_{k_x-1} = a_1$. Therefore, since K is a smallest prism, $f_{j_x+1} = b_1$ and $f_{k_x-1} = a_1$. It follows that $j_x = k_x$ and f_{j_x} is adjacent to both a_1 and b_1, so $f_{j_x} = m_1$, contradicting the fact that K has a tidy frame. □

4.3. *Let x and y be vertices each with at least two nonadjacent neighbors in a path $P = v_1$-$\cdots$-v_n. If there do not exist paths A and B between x and y with $A^* \subset V(P)$ and $B^* \subset V(P)$ such that A^* and B^* are anticomplete to each other, then for some $1 \leq i \leq n$, $N(x) \cap V(P) \subset \{v_1, \ldots, v_{i+1}\}$ and $N(y) \cap V(P) \subset \{v_i, \ldots, v_n\}$.*

Proof. Let v_s and v_t be the neighbors of x in P with s minimum and t maximum. Let v_p and v_r be the neighbors of y in P with p minimum and

r maximum. Since x and y each have two nonadjacent neighbors in P, $s+1<t$ and $p+1<r$. If $p<s$ and $r>t$ then there are paths A and B between x and y with $A^* \subseteq \{v_p, \ldots, v_s\}$ and $B^* \subseteq \{v_t, \ldots, v_r\}$ and A^* and B^* are anticomplete to each other, a contradiction. So either $p \geq s$ or $r \leq t$. By symmetry, we may assume that $p \geq s$. If $r \leq t$ then there are paths A and B between x and y with $A^* \subseteq \{v_s, \ldots, v_p\}$ and $B^* \subseteq \{v_r, \ldots, v_t\}$ and A^* and B^* are anticomplete to each other, a contradiction. Therefore, $r>t$. If $p<t-1$ then there are paths A and B between x and y with $A^* \subseteq \{v_s, \ldots, v_p\}$ and $B^* \subseteq \{v_t, \ldots, v_r\}$ and A^* and B^* are anticomplete to each other, a contradiction. So $p \geq t-1$, and the theorem holds with $i=t-1$. This proves 4.3. □

4.4. *Every K-major vertex has neighbors in at least two of P_1, P_2 and P_3.*

Proof. Let v be a K-major vertex with $N(v) \cap V(K) \subseteq V(P_1)$. Since v is major, $j_v + 2 < k_v$. The subgraph induced by G on $V(K) \cup \{v\} \setminus \{f_{j_v+1}, \ldots, f_{k_v-1}\}$ is a prism that is smaller than K. Therefore, $N(v) \cap V(K) \not\subseteq V(P_1)$. By symmetry between P_1, P_2 and P_3, it follows that $N(v) \cap V(K) \not\subseteq V(P_2)$ and $N(v) \cap V(K) \not\subseteq V(P_3)$. This proves 4.4. □

4.5. *Let x and y be K-major vertices that are not adjacent to each other. Then for some i, j with $1 \leq i < j \leq 3$, $N(x) \cap V(K) \subseteq V(P_i) \cup V(P_j)$ and $N(y) \cap V(K) \subseteq V(P_i) \cup V(P_j)$.*

Proof. We first prove that each of x and y has neighbors in exactly two of P_1, P_2 and P_3.

(1) *None of x and y has neighbors in all of P_1, P_2, P_3.*

Assume that x and y both have neighbors in all three of P_1, P_2 and P_3. Then there are three paths between x and y each with interior in P_i^* for $i \in \{1,2,3\}$. These paths form a theta whose vertex set is contained in $(V(K) \setminus \{a_1, a_2, a_3, b_1, b_2, b_3\}) \cup \{x, y\}$ and is thus smaller than K, a contradiction. Therefore, at least one of x and y does not have neighbors in all three paths P_1, P_2 and P_3.

We may assume by symmetry that $N(y) \cap V(P_1) = \emptyset$. Also, we may assume that x has neighbors in all of P_1, P_2 and P_3, because otherwise (1) holds. By 4.4, y has neighbors in both P_2 and P_3. By 4.1, it has at least two neighbors in one of these paths, say P_2. Then by 4.2, it has two nonadjacent neighbors in P_2.

Since x and y both have neighbors in P_3, there is a path R between them with interior in P_3^*. Let P be the path h_{s_y}-P_2-a_2-a_1-P_1-b_1-b_2-P_2-h_{t_y}. Assume that x has two nonadjacent neighbors in P. Then by 4.3, there are two paths R' and R'' between x and y whose interiors are anticomplete to each other and contained in $V(P)$. Since $V(P)$ is anticomplete to P_3^*, the paths x-R-y, x-R'-y and x-R''-y form a theta, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{a_3, b_3, h_{s_y+1}\}$, K' is smaller than K, a contradiction. Therefore, x does not have two nonadjacent neighbors in P. Since P_1 is a subpath of P, this implies that x does not have two nonadjacent neighbors in P_1, so $j_x \geq k_x - 1$. Also, since $N(x) \cap P_1^* \neq \emptyset$, it follows that $N(x) \cap (V(P) \setminus V(P_1)) = \emptyset$. Thus $s_x > s_y$ and $t_x < t_y$.

If $s_x < t_y - 1$ then the paths x-R-y, x-f_{k_x}-P_1-b_1-b_2-P_2-h_{t_y}-y, and a path between x and y with interior in $\{h_{s_y}, \ldots, h_{s_x}\}$ form a theta, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{h_{t_y-1}, a_3, b_3\}$, K' is smaller than K, a contradiction. Therefore, $s_x \geq t_y - 1$. Since $s_x \leq t_x$ and $t_x < t_y$, this proves that $s_x = t_y - 1 = t_x$. If $t_x > s_y + 1$, then the paths x-R-y, x-f_{j_x}-P_1-a_1-a_2-P_2-h_{s_y}-y and a path between x and y with interior in $\{h_{t_x}, h_{t_y}\}$ form a theta, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{h_{s_y+1}, a_3, b_3\}$, K' is smaller than K, a contradiction. Therefore, $t_x \leq s_y + 1$. Since $t_x = s_x > s_y$, this proves that $s_x = t_x = s_y + 1$. Therefore, $s_y + 2 = t_y$ and $N(x) \cap V(P_2) = \{h_{s_y+1}\}$.

If y is not adjacent to h_{s_y+1} then $|N(y) \cap V(P_2)| = 2$ and $N(y) \cap V(P_1) = \emptyset$, contradicting 4.2. Therefore, y is adjacent to h_{s_y+1}. If $j_x = k_x$ then the subgraph induced by G on $V(P_1) \cup V(P_2) \cup \{x\}$ is a theta smaller than K, a contradiction. Therefore, $j_x < k_x$, and since $j_x \geq k_x - 1$, it follows that $j_x = k_x - 1$. If $c_y < c_x$ then the triangles $\{h_{s_y}, h_{s_y+1}, y\}$ and $\{a_1, a_2, a_3\}$ and the paths h_{s_y}-P_2-a_2, h_{s_y+1}-x-f_{j_x}-P_1-a_1 and y-g_{c_y}-P_3-a_3 form a prism, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{b_1, b_2, b_3\}$, K' is smaller than K, a contradiction. Therefore $c_y \geq c_x$. If $c_x < c_y$ then the triangles $\{x, f_{j_x}, f_{k_x}\}$ and $\{a_1, a_2, a_3\}$ and the paths x-g_{c_x}-P_3-a_3, f_{j_x}-P_1-a_1 and f_{k_x}-P_1-b_1-b_2-P_2-h_{t_y}-y-h_{s_y}-P_2-a_2 form a prism, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{h_{s_y+1}, b_3, g_{c_y}\}$, K' is smaller than K, a contradiction. Therefore, $c_x = c_y$. By the symmetry of the prism K, it follows by the same argument that $d_x = d_y$.

If $c_x = d_x$ then the induced subgraph with vertex set $V(P_2) \cup V(P_3) \cup \{x\}$ is a theta smaller than K, a contradiction. So $c_x < d_x$, which means that $c_y < d_y$. If $c_y < d_y - 1$ then the paths x-g_{c_y}-y, x-g_{d_y}-y and x-h_{s_y+1}-y form a theta smaller than K, a contradiction. Therefore, $c_y = d_y - 1$. But then $|N(y) \cap V(P_3)| = 2$ and $N(y) \cap V(P_1) = \emptyset$, which contradicts 4.2. This proves (1).

By 4.4 and (1), each of x and y has neighbors in precisely two of the paths P_1, P_2 and P_3. Assume that $N(x) \cap V(K)$ and $N(y) \cap V(K)$ are not both contained in the union of the same two of $V(P_1)$, $V(P_2)$ and $V(P_3)$. By symmetry, we may assume that $N(x) \cap V(P_3) = \emptyset$ and that $N(y) \cap V(P_1) = \emptyset$.

Assume that $j_x < k_x - 1$ and $c_y < d_y - 1$; this means that x has two nonadjacent neighbors in P_1 and y has two nonadjacent neighbors in P_3. Since both x and y have neighbors in P_2, there exists a path R between x and y with interior in P_2^*. Then the three paths x-f_{j_x}-P_1-a_1-a_3-P_3-g_{c_y}-y, x-f_{k_x}-P_1-b_1-b_3-P_3-g_{d_y}-y and x-R-y form a theta, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{a_2, b_2, f_{j_x+1}\}$, K' is smaller than K, a contradiction. This implies that either x does not have two nonadjacent neighbors in P_1 or y does not have two nonadjacent neighbors in P_3. By symmetry we may assume that x does not have two nonadjacent neighbors in P_1, so $j_x \geq k_x - 1$. By 4.2, $N(x) \cap V(P_1)$ does not consist of two adjacent vertices. Therefore, $|N(x) \cap V(P_1)| = 1$, and $j_x = k_x$.

Since x is K-major, by 4.1, $|N(x) \cap V(P_2)| \geq 2$ and by 4.2, $N(x) \cap V(P_2)$ does not consist of two adjacent vertices. Therefore, x has two nonadjacent neighbors in P_2, and $s_x < t_x - 1$.

First assume that $|N(y) \cap V(P_2)| = 1$. Then $s_y = t_y$, and by 4.2 and 4.1, y has two nonadjacent neighbors in P_3, so $c_y < d_y - 1$. If $s_y < t_x - 1$ then the paths y-g_{d_y}-P_3-b_3-b_2-P_2-h_{t_x}-x, y-g_{c_y}-P_3-a_3-a_1-P_1-f_{j_x}-x, and a path between x and y with interior in $V(h_{s_y}$-P_2-$h_{s_x})$ form a theta, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{h_{t_x-1}, a_2, b_1\}$, K' is smaller than K, a contradiction. Therefore, $s_y \geq t_x - 1$. From the symmetry, $s_y \leq s_x + 1$. So we have that $t_x - 1 \leq s_y \leq s_x + 1$ and $s_x < t_x - 1$, from which it follows that $s_x + 1 = s_y = t_x - 1$. If x is not adjacent to h_{s_x+1} then $|N(x) \cap V(P_2)| = 2$ and $N(x) \cap V(P_3) = \emptyset$, contradicting 4.2. Therefore, x is adjacent to h_{s_x+1}. The triangles $\{h_{s_x}, h_{s_x+1}, x\}$ and $\{a_1, a_2, a_3\}$ and the paths h_{s_x}-P_2-a_2, h_{s_x+1}-y-g_{c_y}-P_3-a_3 and x-f_{j_x}-P_1-a_1 form a prism, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{b_1, b_2, b_3\}$, K' is smaller than K, a contradiction. This proves that $|N(y) \cap V(P_2)| > 1$.

By 4.2, y has more than two neighbors in P_2, and $s_y < t_y - 1$. Let Q be the path x-f_{j_x}-P_1-a_1-a_3-P_3-g_{c_y}-y. If there are two paths A and B between x and y whose interiors are anticomplete to each other and contained in P_2^*, then the paths Q, A and B form a theta, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{b_1, b_2, b_3\}$, K' is smaller than K, a contradiction. Therefore, there are no two paths between x and y whose interiors are anticomplete to each other and contained in P_2^*. Then by 4.3, either $N(y) \cap V(P_2) \subseteq \{h_1, \dots, h_{s_x+1}\}$ or $N(y) \cap V(P_2) \subseteq \{h_{t_x-1}, \dots, h_m\}$. By symmetry, we may assume that $N(y) \cap V(P_2) \subseteq \{h_1, \dots, h_{s_x+1}\}$, which means that $t_y \leq s_x + 1$.

If $t_y \leq s_x$, then the paths y-h_{s_y}-P_2-a_2-a_1-P_1-f_{j_x}-x, y-g_{d_y}-P_3-b_3-b_2-P_2-h_{t_x}-x, and a path between x and y with interior in $V(h_{t_y}$-P_2-$h_{s_x})$ form a theta,

K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{a_3, b_1, h_{s_x+1}\}$, K' is smaller than K, a contradiction. Therefore, $t_y > s_x$, and since $t_y \leq s_x + 1$, it follows that $t_y = s_x + 1$ and y is adjacent to h_{s_x+1}.

Assume that $s_y < s_x - 1$ and $t_x > t_y + 1$. Then the paths x-f_{j_x}-P_1-a_1-a_2-P_2-h_{s_y}-y, x-h_{t_x}-P_2-b_2-b_3-P_3-g_{d_y}-y and a path between x and y with interior in $\{h_{s_x}, h_{s_x+1}\}$ form a theta, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{a_3, h_{s_x-1}, h_{t_x-1}\}$, K' is smaller than K, a contradiction. Therefore, either $s_y \geq s_x - 1$ or $t_x \leq t_y + 1$.

First assume that $s_y \geq s_x - 1$. Then since $s_x = t_y - 1$ and $s_y < t_y - 1$, this implies that $s_y = s_x - 1$. If y is adjacent to h_{s_x} then the triangles $\{h_{s_x-1}, h_{s_x}, y\}$ and $\{a_1, a_2, a_3\}$ and the paths h_{s_x-1}-P_2-a_2, h_{s_x}-x-f_{j_x}-P_1-a_1 and y-g_{c_y}-P_3-a_3 form a prism, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{b_1, b_2, b_3\}$, K' is smaller than K, a contradiction. Therefore, y is not adjacent to h_{s_x}. Then since $t_y = s_y + 2$, this means that $|N(y) \cap V(P_2)| = 2$, which contradicts 4.2. This proves that $s_y < s_x - 1$, so we may now assume that $t_x \leq t_y + 1$. Since $t_y = s_x + 1$ and $t_x > s_x + 1$, this implies that $t_x = s_x + 2$. If x is adjacent to h_{s_x+1} then the triangles $\{h_{s_x+1}, h_{t_x}, x\}$ and $\{b_1, b_2, b_3\}$ and the paths h_{s_x+1}-y-g_{d_y}-P_3-b_3, h_{t_x}-P_2-b_2, and x-f_{k_x}-P_1-b_1 form a prism, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{a_1, a_2, a_3\}$, K' is smaller than K, a contradiction. Therefore, x is not adjacent to h_{s_x+1}. Then since $t_x = s_x + 2$, this means that $|N(x) \cap V(P_2)| = 2$, which contradicts 4.2. This proves 4.5. □

4.6. *Two K-major vertices that are not adjacent have a common neighbor in K.*

Proof. Let x and y be K-major vertices that are not adjacent to each other. Assume that $N(x) \cap V(K) \cap N(y) = \emptyset$. From 4.5 we may assume that $(N(x) \cup N(y)) \cap V(K) \subset P_1^* \cup P_2^*$. By 4.4, each of P_1^* and P_2^* contains a neighbor of x and a neighbor of y. Let C be the cycle induced by G on $V(P_1) \cup V(P_2)$. A *gap* is a minimal path in C containing a neighbor of x and a neighbor of y. Every edge of C is in at most one gap and any vertex common to two distinct gaps is an end of both. Since x and y have no common neighbor in C, there is an even number of gaps. For a gap P, let P^+ be the path between x and y whose interior is $V(P)$. Let $R_1, \dots, R_k$ be the gaps of C numbered in their order on C, such that the vertex set of the path of C sharing one end with R_i and the other end with R_{i+1}, and with interior disjoint from all the gaps, is anticomplete to y if i is odd and anticomplete to x if i is even. Call a gap a *small gap* if it consists of one edge and shares both of its ends with other gaps.

(1) *There are at most six gaps.*

Suppose there are at least eight gaps. Assume that for some i, R_i is not a small gap. Then the paths R_{i+1}^+, R_{i+4}^+ and R_{i+7}^+ form a theta, K'. It is smaller than K because each of R_i, $R_{i+2} \cup R_{i+3}$ and $R_{i+5} \cup R_{i+6}$ contains at least one vertex not in $V(K')$. Therefore, every gap is small. Then every edge of C is in a gap and every gap consists of one edge, so $V(C) \subseteq N(x) \cup N(y)$. But this contradicts the fact that x and y have no neighbors in $\{a_1, a_2, b_1, b_2\} \subseteq C$. This proves (1).

(2) *There are at most four gaps.*

Suppose there are six gaps. Assume that no gap is small. The paths R_1^+, R_3^+ and R_5^+ form a theta. It is smaller than K because R_2, R_4 and R_6 each contain at least one vertex not in the theta since they are not small gaps. So there is a small gap; by symmetry let it be R_1. Assume next that none of R_2, R_4 or R_6 is a small gap. Then the paths R_1^+, R_3^+ and R_5^+ form a theta smaller than K. Therefore one of R_2, R_4 or R_6 is small.

First assume that R_2 is a small gap. Then $R_1 \cup R_2$ is a three-vertex path whose vertices are contained in $N(x) \cup N(y)$, and so it is a subpath of either P_1 or P_2 in K, because x and y have no neighbors in $\{a_1, a_2, b_1, b_2\}$. By symmetry assume it is a subpath of P_1. Both x and y have a neighbor in P_2, so there is a path Q between them with interior in P_2. Let u be the vertex that is an end of both R_1 and R_2. The three paths u-R_1^+-y, u-R_2^+-y and u-x-Q-y form a theta, K'. Since $V(K')$ is disjoint from $\{a_1, a_2, a_3, b_1, b_2, b_3\}$, it follows that the theta K' is smaller than K, a contradiction. This proves that R_2 is not a small gap. From the symmetry, we deduce that R_6 is not a small gap either.

Therefore R_4 is a small gap. From the argument in the previous paragraph with R_4 playing the role of R_1, we deduce that neither of R_3, R_5 is a small gap. Since R_1 is a small gap and R_5 is not, we deduce that $V(R_1) \cap V(R_6) \neq \emptyset$, and that $V(R_6)$ is anticomplete to $V(R_4)$. Let $V(R_1) \cap V(R_6) = \{v\}$. The paths x-R_4^+-y-v, $R_6^+ \setminus y$, and $R_1^+ \setminus y$ form a theta smaller than K, a contradiction.

This proves (2).

(3) *There are at most two gaps.*

Assume there are four gaps. Let $Q_1, \ldots, Q_4$ be subpaths of C such that none of their interiors meet any gap, and such that for $1 \leq i \leq 3$, one end of Q_i is an end of R_i and the other is an end of R_{i+1}, while the ends of Q_4 are ends of R_1 and R_4. The neighbors of x in C are contained in $V(Q_1) \cup V(Q_3)$ and the neighbors of y in C are contained in $V(Q_2) \cup V(Q_4)$.

Assume that Q_1 and Q_3 both have length at most one. Then since x has neighbors in both P_1 and P_2 but no neighbors in $\{a_1, a_2, b_1, b_2\}$, one of Q_1 and Q_3 is a subpath of P_1 and the other is a subpath of P_2. But then x has at most two neighbors in each of P_1 and P_2, contradicting 4.1 or 4.2. This shows that at least one of Q_1 and Q_3 has length at least two. By symmetry, we may assume that Q_1 does. Therefore $V(R_1)$ is anticomplete to $V(R_2)$.

By the same argument applied to Q_2 and Q_4, we deduce that at least one of Q_2 and Q_4 has length at least two. From the symmetry, we may assume that Q_2 has length at least two, and therefore $V(R_2)$ is anticomplete to $V(R_3)$. Since the paths R_1^+, R_2^+ and R_3^+ do not form a theta smaller than K, it follows that $|V(Q_3)| = |V(Q_4)| = 1$ and R_4 is a small gap. Let $V(Q_3) = \{q_3\}$ and $V(Q_4) = \{q_4\}$. Since $V(R_2)$ is anticomplete to both $V(R_1)$ and $V(R_3)$, the paths x-q_3-q_4, $R_1^+ \setminus y$, and x-R_2^+-y-q_4 form a theta smaller that K, a contradiction. This proves (3).

Since x and y each have a neighbor in C, there are at least two gaps; so from (3) there are two gaps. Since x and y each have neighbors in both P_1 and P_2, there is a gap contained in P_1 and a gap contained in P_2. The subgraph $C \setminus (R_1 \cup R_2)$ is the disjoint union of two paths, Q_1 and Q_2, such that Q_1 contains no neighbors of x and Q_2 contains no neighbors of y. By 4.1 and 4.2, each of x and y has two nonadjacent neighbors in one of P_1 and P_2, so each of x and y has a neighbor in C that is anticomplete to $V(R_1) \cup V(R_2)$. Therefore, there is a path S between x and y whose interior is anticomplete to $V(R_1) \cup V(R_2)$ and is contained in $V(K) \setminus (V(R_1) \cup V(R_2))$. The paths S, R_1^+ and R_2^+ form a theta smaller than K, a contradiction. This proves 4.6. □

A *broom* is a vertex $v \in V(G)$ such that for some $\{i_1, i_2, i_3\} = \{1, 2, 3\}$,

- $N(v) \cap V(P_{i_1}) = \emptyset$,
- $|N(v) \cap V(P_{i_2})| = 1$, and
- $N(v) \cap V(P_{i_3})$ is the vertex set of a three-vertex path.

4.7. *If x and y are two K-major vertices that are not adjacent to each other, then one of them is a broom, and for some i, j such that $1 \leq i < j \leq 3$, $N(x) \cap V(P_i) \cap N(y) \neq \emptyset$ and $N(x) \cap V(P_j) \cap N(y) \neq \emptyset$.*

Proof. By 4.5, we may assume that $N(x) \cap V(K) \subseteq P_1^* \cup P_2^*$ and $N(y) \cap V(K) \subseteq P_1^* \cup P_2^*$. By 4.4, x and y each have a neighbor in both P_1^* and in P_2^*. Also, by 4.6, x and y have a common neighbor in K, and by the symmetry between P_1 and P_2 we may assume that $f_r \in P_1^*$ is a neighbor of both x and y.

Let A be the path induced by G on $V(P_1) \cup V(P_2) \setminus \{f_{r-1}, f_r, f_{r+1}\}$.

(1) *x and y do not both have a neighbor in $V(A) \cap V(P_1)$.*

Suppose both x and y have a neighbor in $V(A) \cap V(P_1)$. Then there is a path, S, between x and y with interior in $(V(A) \cap V(P_1)) \cup V(P_3)$. Since x and y each have a neighbor in P_2^* there is a path, R, between them with interior in P_2^*. The paths S, x-f_r-y and R form a theta, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{f_{r-1}, f_{r+1}, a_2, b_2\}$, K' is smaller than K, a contradiction. This proves (1).

(2) *If exactly one of x and y has a neighbor in $V(A) \cap V(P_1)$, then the theorem holds.*

Assume that exactly one of x and y has neighbors in $V(A) \cap V(P_1)$. By symmetry between x and y we may assume that $N(x) \cap V(A) \cap V(P_1) \neq \emptyset$ and $N(y) \cap V(A) \cap V(P_1) = \emptyset$. By 4.1 and 4.2, y has two nonadjacent neighbors in at least one of P_1 and P_2. Since y has no neighbors in $V(A) \cap V(P_1)$, this means that it either has two nonadjacent neighbors in P_2, so $t_y > s_y + 1$, or it is adjacent to both f_{r-1} and f_{r+1}.

First assume that $t_y > s_y + 1$. Since $N(x) \cap V(A) \cap V(P_1) \neq \emptyset$ and $N(x) \cap V(P_2) \neq \emptyset$, each of x and y has two nonadjacent neighbors in A. If there are two paths S and R between x and y whose interiors are anticomplete to each other and are contained in $V(A)$, then the paths x-f_r-y, S and R form a theta smaller than K, a contradiction. Therefore, by 4.3 we may assume that $N(x) \cap V(A) \subseteq V(a_1\text{-}P_1\text{-}f_{r-2}) \cup V(a_2\text{-}P_2\text{-}h_{s_y+1})$, which means that $t_x \leq s_y + 1$. Since $N(x) \cap V(A) \cap V(P_1) \neq \emptyset$, $j_x < r - 1$. If $s_x < t_y - 1$, then the paths x-f_r-y, x-f_{j_x}-P_1-a_1-a_3-P_3-b_3-b_2-P_2-h_{t_y}-y and a path between x and y with interior in $V(h_{s_x}\text{-}P_2\text{-}h_{s_y})$ form a theta, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{a_2, b_1, f_{r-1}\}$, K' is smaller than K, a contradiction. Therefore, $s_x \geq t_y - 1$, and since $t_x \leq s_y + 1$ and $t_y > s_y + 1$, it follows that $s_y + 1 = t_y - 1$ and $N(x) \cap V(P_2) = \{h_{s_y+1}\}$. By 4.2, $|N(y) \cap V(P_2)| \neq 2$, so y is adjacent to h_{s_y+1}. Also by 4.2, $|N(y) \cap V(P_1)| \neq 2$, so y is either complete or anticomplete to $\{f_{r-1}, f_{r+1}\}$.

Suppose that y is $\{f_{r-1}, f_{r+1}\}$-complete. If x is not adjacent to f_{r+1} then the triangles $\{f_r, f_{r+1}, y\}$ and $\{b_1, b_2, b_3\}$ and the paths f_r-x-f_{j_x}-P_1-a_1-a_3-P_3-b_3, f_{r+1}-P_1-b_1 and y-h_{t_y}-P_2-b_2 form a prism, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{h_{s_y}, h_{s_y+1}, f_{r-1}\}$, K' is smaller than K, a contradiction. So x is adjacent to f_{r+1}, and the paths x-f_{r+1}-y, x-h_{s_y+1}-y and a path between x and y with interior in $V(f_{j_x}\text{-}P_1\text{-}f_{r-1})$ form a theta smaller than K, a contradiction. Therefore, y is anticomplete to $\{f_{r-1}, f_{r+1}\}$, and so y is a broom and the theorem holds.

Since we have shown that the theorem holds when $t_y > s_y + 1$, we may now assume that $t_y \leq s_y + 1$ and that y is $\{f_{r-1}, f_{r+1}\}$-complete. By 4.2, $|N(y) \cap V(P_2)| \neq 2$, so $t_y = s_y$. Since $N(x) \cap V(A) \cap V(P_1) \neq \emptyset$, either $j_x < r - 1$ or $k_x > r + 1$. By symmetry we may assume that $j_x < r - 1$.

Suppose that $|N(x) \cap V(P_2)| = 1$. If x is adjacent to f_{r+1} then the path $x\text{-}f_{r+1}\text{-}y$, a path between x and y with interior in $V(f_{j_x}\text{-}P_1\text{-}f_{r-1})$ and a path between x and y with interior in P_2^* form a theta smaller than K, a contradiction. Therefore, x is not adjacent to f_{r+1}. If $s_x < s_y$ then the triangles $\{f_r, f_{r+1}, y\}$ and $\{b_1, b_2, b_3\}$ and the paths $f_r\text{-}x\text{-}f_{j_x}\text{-}P_1\text{-}a_1\text{-}a_3\text{-}P_3\text{-}b_3$, $f_{r+1}\text{-}P_1\text{-}b_1$ and $y\text{-}h_{s_y}\text{-}P_2\text{-}b_2$ form a prism, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{f_{r-1}, a_2, h_{s_x}\}$, K' is smaller than K, a contradiction. If $s_x > s_y$, then the triangles $\{f_r, f_{r+1}, y\}$ and $\{a_1, a_2, a_3\}$ and the paths $f_r\text{-}x\text{-}f_{j_x}\text{-}P_1\text{-}a_1$, $f_{r+1}\text{-}P_1\text{-}b_1\text{-}b_3\text{-}P_3\text{-}a_3$ and $y\text{-}h_{s_y}\text{-}P_2\text{-}a_2$ form a prism, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{f_{r-1}, b_2, h_{s_x}\}$, K' is smaller than K, a contradiction. This proves that $s_x = s_y$, and y is a broom, so the theorem holds. So we may assume that $|N(x) \cap V(P_2)| > 1$, and it follows from 4.2 that $s_x < t_x - 1$.

Let Q be the path between x and y with interior in $V(f_{j_x}\text{-}P_1\text{-}f_{r-1})$. If $t_x \leq s_y$ then the paths Q, $x\text{-}h_{t_x}\text{-}P_2\text{-}h_{s_y}\text{-}y$ and $x\text{-}h_{s_x}\text{-}P_2\text{-}a_2\text{-}a_3\text{-}P_3\text{-}b_3\text{-}b_1\text{-}P_1\text{-}f_{r+1}\text{-}y$ form a theta, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{f_r, a_1, b_2\}$, K' is smaller than K, a contradiction. Therefore, $t_x > s_y$. If $s_x \geq s_y$ then the paths Q, $x\text{-}h_{s_x}\text{-}P_1\text{-}h_{s_y}\text{-}y$ and $x\text{-}h_{t_x}\text{-}P_2\text{-}b_2\text{-}b_1\text{-}P_1\text{-}f_{r+1}\text{-}y$ form a theta, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{a_1, a_2, a_3\}$, K' is smaller than K, a contradiction. Therefore, $s_x < s_y$. If x is not adjacent to h_{s_y}, then the path $h_{s_y}\text{-}y\text{-}f_r\text{-}x$, a path between h_{s_y} and x with interior in $V(h_{s_y}\text{-}P_2\text{-}h_{s_x})$ and a path between h_{s_y} and x with interior in $V(h_{s_y}\text{-}P_2\text{-}h_{t_x})$ form a theta smaller than K, a contradiction. Thus x is adjacent to h_{s_y}, and y is a broom, so the theorem holds. This proves (2).

(3) *If neither x nor y has a neighbor in $V(A) \cap V(P_1)$, then the theorem holds.*

Suppose that neither x nor y has a neighbor in $V(A) \cap V(P_1)$. By 4.2, each of x and y is either complete or anticomplete to $\{f_{r-1}, f_{r+1}\}$. If both x and y are $\{f_{r-1}, f_{r+1}\}$-complete, then the paths $x\text{-}f_{r-1}\text{-}y$, $x\text{-}f_{r+1}\text{-}y$ and a path between x and y with interior in P_2^* form a theta smaller than K, a contradiction. This means that at most one of x and y is $\{f_{r-1}, f_{r+1}\}$-complete.

Assume that exactly one of x and y is $\{f_{r-1}, f_{r+1}\}$-complete. By symmetry, we may assume that x is, and that y is anticomplete to $\{f_{r-1}, f_{r+1}\}$. By 4.1 and 4.2, $|N(y) \cap V(P_2)| \geq 3$, so $s_y + 1 < t_y$.

Suppose that $|N(x) \cap V(P_2)| > 1$; then by 4.2, $|N(x) \cap V(P_2)| \geq 3$ so $s_x + 1 < t_x$. If there are two paths S and R between x and y whose interiors are anticomplete to each other and contained in P_2^*, then S, R and x-f_r-y form a theta smaller than K, a contradiction. So from 4.3, we may assume that $t_x \leq s_y + 1$. From this it follows that $s_x < s_y$. Then the triangles $\{f_r, f_{r+1}, x\}$ and $\{b_1, b_2, b_3\}$ and the paths f_r-y-h_{t_y}-P_2-b_2, f_{r+1}-P_1-b_1 and x-h_{s_x}-P_2-a_2-a_3-P_3-b_3 form a prism, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{a_1, h_{s_y}, f_{r-1}\}$, K' is smaller than K, a contradiction. This proves that $|N(x) \cap V(P_2)| = 1$.

If $s_x < s_y$ then the triangles $\{f_{r-1}, f_r, x\}$ and $\{a_1, a_2, a_3\}$ and the paths f_{r-1}-P_1-a_1, f_r-y-h_{t_y}-P_2-b_2-b_3-P_3-a_3 and x-h_{s_x}-P_2-a_2 form a prism, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{h_{s_y}, b_1, f_{r+1}\}$, K' is smaller than K, a contradiction. Therefore $s_x \geq s_y$, and from the symmetry $s_x \leq t_y$.

If y is not adjacent to h_{s_x}, then $s_y < s_x$ and $t_y > s_x$. The path y-f_r-x-h_{s_x} and two paths between y and h_{s_x} with interiors in $V(h_{s_y}$-P_2-$h_{s_x})$ and $V(h_{t_y}$-P_2-$h_{s_x})$ form a theta smaller than K, a contradiction. So y is adjacent to h_{s_x}. Since x is a broom, the theorem holds.

Since we have shown that the theorem holds when one of x and y is $\{f_{r-1}, f_{r+1}\}$-complete, we may now assume that $\{x, y\}$ is anticomplete to $\{f_{r-1}, f_{r+1}\}$. Now $N(x) \cap V(P_1) = N(y) \cap V(P_1) = \{f_r\}$, so by 4.1 and 4.2, x and y each have two nonadjacent neighbors in P_2. If there are two paths S and R between x and y whose interiors are anticomplete to each other and contained in P_2^*, then the paths S, R and x-f_r-y form a theta smaller than K, a contradiction. Therefore, by 4.3, we may assume that $t_x \leq s_y + 1$.

If $t_x \leq s_y$ then the paths x-f_r-y, x-h_{t_x}-P_2-h_{s_y}-y and x-h_{s_x}-P_2-a_2-a_3-P_3-b_3-b_2-P_2-h_{t_y}-y form a theta, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{a_1, b_1, h_{s_y+1}\}$, K' is smaller than K, a contradiction. Therefore $t_x > s_y$ so $t_x = s_y + 1$.

If $s_x < s_y - 1$ and $t_y > t_x + 1$ then the paths x-f_r-y, x-h_{s_x}-P_2-a_2-a_3-P_3-b_3-b_2-P_2-h_{t_y}-y and a path between x and y with interior in $\{h_{s_y}, h_{t_x}\}$ form a theta, K'. Since $V(K') \subseteq V(K) \cup \{x, y\} \setminus \{a_1, b_1, h_{s_y-1}\}$, K' is smaller than K, a contradiction. Therefore, either $s_x \geq s_y - 1$ or $t_y \leq t_x + 1$. From the symmetry we may assume that $s_x \geq s_y - 1$. Then since $s_x < t_x - 1$ and $t_x = s_y + 1$, it follows that $s_x = s_y - 1$. If x is not adjacent to h_{s_y} then $|N(x) \cap V(P_2)| = 2$, contradicting 4.2. So x is adjacent to h_{s_y}. Then x is a broom and the theorem holds. This proves (3).

Now 4.7 follows from (1), (2) and (3). □

Let M be the subgraph of G induced on the set of K-major vertices. The following algorithm constructs a superset of $V(M)$ that is disjoint from $V(K)$ and is used for cleaning.

4.8. *There is an algorithm with the following specifications:*

- **Input:** *A graph G.*
- **Output:** *A sequence of subsets $X_1, \dots, X_r$ of $V(G)$, with $r \leq 2|V(G)|^6$, such that for every smallest prism K in G, if K has a tidy frame then one of $X_1, \dots, X_r$ is disjoint from $V(K)$ and contains all K-major vertices.*
- **Running Time:** *$O(|V(G)|^7)$.*

Proof. The algorithm is as follows. For each triple of vertices (a, b, c) compute the set U of vertices complete to $N(a, b) \cup \{c\}$. List all subsets $W \subseteq V(G)$ of at most three vertices. For each pair of (a, b, c) and W, compute the subset

$$N(a, b) \cup \{c\} \cup U \setminus W.$$

Label the subsets generated $Y_1, \dots, Y_s$, and label the list of subsets $N(a, b)$ by $N_1, \dots, N_p$. Enumerate all quintuples of vertices (u, w, z_1, z_2, z_3). For each, compute the subset

$$Z(u, w, z_1, z_2, z_3) = \{u\} \cup (N(u) \setminus \{w, z_1, z_2, z_3\}) \cup N(w, z_1) \cup N(w, z_2) \cup N(w, z_3).$$

Label the subsets generated $Z_1, \dots, Z_t$. Output the subsets

$$\emptyset, N_1, \dots, N_p, Y_1, \dots, Y_s, Z_1, \dots, Z_t.$$

This completes the description of the algorithm.

Since $p \leq |V(G)|^2$, $s \leq |V(G)|^6$ and $t \leq |V(G)|^5$, the number of subsets in the output sequence is $r \leq 2|V(G)|^6$. For each triple (a, b, c), the subset U can be computed in quadratic time and the subsets W can be enumerated in cubic time. Computing $N(a, b) \cup \{c\} \cup U \setminus W$ takes linear time, so the time taken to generate the sequence $N_1, \dots, N_p, Y_1, \dots, Y_s$ is $O(|V(G)|^7)$. Since each of the $|V(G)|^5$ subsets Z can be computed in linear time, their computation does not affect the total running time.

Let K be a smallest prism in G, and assume K has a tidy frame. Let M be the subgraph induced by G on the set of K-major vertices. Suppose M is the complete graph. We may assume that there exists a set $\{a, b\} \subseteq V(K)$ that is not a subset of the vertex set of a 3-vertex subpath of K, because if no such choice is possible then $V(M) = \emptyset$, which is in the output. We may also assume that, subject to these conditions, $\{a,b\}$ is chosen with $|N(a,b)|$ maximum. Then $N(a,b) \subseteq V(M)$. If $N(a,b) = V(M)$ then $V(M)$ is in the output sequence. Otherwise there is a vertex $c \in V(M) \setminus N(a,b)$. Then since $N(a,b) \cup \{c\} \subseteq V(M)$ and M is a complete graph, the set U of vertices complete to $N(a,b) \cup \{c\}$ contains $V(M) \setminus (N(a,b) \cup \{c\})$; so

$$V(M) \subseteq N(a,b) \cup \{c\} \cup U.$$

If $U \cap V(K)$ is not a subset of the vertex set of a 3-vertex path of K, then there are two vertices $a', b' \in U \cap V(K)$ such that $\{a', b'\}$ is not a subset of the vertex set of a 3-vertex path of K. Since $N(a,b) \cup \{c\} \subseteq N(a', b')$, this contradicts the choice of a and b. So $|U \cap V(K)| \leq 3$ and we can choose $W = U \cap V(K)$. Then the subset $N(a,b) \cup \{c\} \cup U \setminus W$ contains $V(M)$ and since $(N(a,b) \cup \{c\}) \cap V(K) = \emptyset$, $N(a,b) \cup \{c\} \cup U \setminus W$ is disjoint from $V(K)$. The subset in the output corresponding to this choice of (a,b,c) and W contains $V(M)$ and is disjoint from $V(K)$.

Therefore, we may assume that M is not a complete graph. Then there exist two vertices in $V(M)$ that are not adjacent to each other. By 4.7, one of these is a broom. So we may choose u to be a broom such that z_1, z_2, z_3 and w are its neighbors in K and z_1-z_2-z_3 is a path of K. By 4.7, any nonneighbor of u in $V(M)$ is contained in the subset $N(w, z_1) \cup N(w, z_2) \cup N(w, z_3)$. Therefore the subset $Z(u, w, z_1, z_2, z_3)$ contains $V(M)$. Since w is in the interior of a different path of the prism K from that containing $\{z_1, z_2, z_3\}$, the subset $N(w, z_1) \cup N(w, z_2) \cup N(w, z_3)$ is disjoint from $V(K)$. Since $N(u) \cap V(K) = \{w, z_1, z_2, z_3\}$, $Z(u, w, z_1, z_2, z_3)$ is disjoint from $V(K)$. This proves 4.8. □

5 The Complete Algorithm

5.1. *There is an algorithm with the following specifications:*

- **Input:** *A graph G.*
- **Output:** *Either:*
  - *a theta in G, or*
  - *a determination that there is no smallest theta in G.*
- **Running Time:** $O(|V(G)|^{22})$.

Proof. Here is the algorithm. Enumerate all 11-tuples $(a, b, s_1, s_2, s_3, m_1, m_2, m_3, t_1, t_2, t_3)$ of vertices of G such that

- a, b, s_1, s_2 and s_3 are all distinct,

- a, b, t_1, t_2 and t_3 are all distinct,
- a is not adjacent to b, and
- a is complete to $\{s_1, s_2, s_3\}$ and b is complete to $\{t_1, t_2, t_3\}$.

We can find all such 11-tuples in time $O(|V(G)|^{11})$. For each 11-tuple do the following.

Remove $N(a) \setminus \{s_1, s_2, s_3\}$ and $N(b) \setminus \{t_1, t_2, t_3\}$ from $V(G)$, and for $i = 1, 2, 3$, if m_i is adjacent to a and b then remove $N(m_i) \setminus \{a, b\}$ from $V(G)$. This can be done in linear time, and after this step every smallest theta in G with the 11-tuple as its frame has a tidy frame. Run the algorithm of 2.2, which takes $O(|V(G)|^4)$ time. Let $X_1, \ldots, X_t$ be the subsets in the output; $t \leq |V(G)|^3$. For every triple of subsets $(X_{i_1}, X_{i_2}, X_{i_3})$ from this list, do the following.

Let G_1 be the graph induced by G on $V(G) \setminus X_{i_1}$. Find shortest paths S'_1 between s_1 and m_1 and T'_1 between m_1 and t_1 in G_1. Next, let G_2 be the graph induced by G_1 on $V(G_1) \setminus X_{i_2}$. Find shortest paths S'_2 between s_2 and m_2 and T'_2 between m_2 and t_2 in G_2. Finally, let G_3 be the graph induced by G_2 on $V(G_2) \setminus X_{i_3}$. Find shortest paths S'_3 between s_3 and m_3 and T'_3 between m_3 and t_3 in G_3. Finding these paths takes quadratic time. Finally, test whether the following are true in the graph G:

- For $1 \leq i \leq 3$, $V(S'_i) \setminus \{m_i\}$ and $V(T'_i) \setminus \{m_i\}$ are disjoint and anticomplete to each other, and
- For $1 \leq i, j \leq 3$ with $i \neq j$, $V(S'_i) \cup V(T'_i)$ and $V(S'_j) \cup V(T'_j)$ are disjoint and anticomplete to each other.

If these conditions are satisfied, output that $\{a, b\}$ and the three paths a-s_1-S'_1-m_1-T'_1-t_1-b, a-s_2-S'_2-m_2-T'_2-t_2-b and a-s_3-S'_3-m_3-T'_3-t_3-b form a theta and stop. Testing these conditions takes quadratic time. So the time it takes to examine one 11-tuple is $O(|V(G)|^{11})$.

After examining all 11-tuples, output that G contains no smallest theta. The total running time is $O(|V(G)|^{22})$. Now we need to prove that this algorithm is correct. Suppose there is a smallest theta K in G; let its half-paths be S_1, T_1, S_2, T_2, S_3 and T_3. Some 11-tuple chosen is the frame of K, and after the first step of the algorithm, K and every smallest theta in G with the same frame as K have a tidy frame.

By 2.2, we may choose X_{i_1} such that the graph G_1 contains K and there are no bad shortcuts across S_1 or T_1 in G_1. Therefore, the subgraph induced by G_1 on $V(K) \setminus (V(S_1) \cup V(T_1)) \cup V(S'_1) \cup V(T'_1)$ is a theta, K_1, in G_1. Since S'_1 and T'_1 are shortest paths, $|V(K_1)| \leq |V(K)|$, so K_1 is also a smallest theta in G_1. Since K_1 is not smaller than K in G it follows that $V(S'_1)\setminus\{m_1\}$ and $V(T'_1) \setminus \{m_1\}$ are anticomplete to each other. We observe that K_1 has the same frame as K, so it has a tidy frame.

Therefore, by 2.2, for some subset X_{i_2} the graph G_2 contains K_1 and all shortcuts across S_2 or T_2 in G_2 are good shortcuts across K_1. Then the subgraph induced by G_2 on $V(K_1) \setminus (V(S_2) \cup V(T_2)) \cup V(S'_2) \cup V(T'_2)$ is a theta, K_2, in G_2. Since S'_2 and T'_2 are shortest paths, $|V(K_2)| \leq |V(K_1)|$, so K_2 is also a smallest theta in G_2. Since K_2 is not smaller than K_1 in G it follows that $V(S'_2) \setminus \{m_2\}$ and $V(T'_2) \setminus \{m_2\}$ are anticomplete to each other. We observe that K_2 has the same frame as K, so it has a tidy frame.

Finally, by 2.2 again, for some subset X_{i_3} the graph G_3 contains K_2 and all shortcuts across S_3 or T_3 in G_3 are good shortcuts across K_2. Then the subgraph induced by G_3 on $V(K_2) \setminus (V(S_3) \cup V(T_3)) \cup V(S'_3) \cup V(T'_3)$ is a theta, K_3, in G_3. Since K_3 is not smaller than K_2 in G it follows that $V(S'_3) \setminus \{m_3\}$ and $V(T'_3) \setminus \{m_3\}$ are anticomplete to each other. We observe that K_3 is the theta that is output by the algorithm when it considers $X_{i_1}, X_{i_2}, X_{i_3}$. This proves that the output of the algorithm is a theta.

Conversely, if the algorithm outputs that there is a theta, then the properties of the 11-tuple and the conditions on the paths that the algorithm checks ensure that this output is actually a theta. □

5.2. *There is an algorithm with the following specifications:*

- **Input:** *A graph G.*
- **Output:** *Either:*
  - a prism in G, or
  - a determination that there is no smallest prism in G.
- **Running Time:** $O(|V(G)|^{35})$.

Proof. Here is the algorithm. Enumerate all 15-tuples of vertices of G

$$(a_1, a_2, a_3, m_1, m_2, m_3, b_1, b_2, b_3, a'_1, a'_2, a'_3, b'_1, b'_2, b'_3)$$

such that

- $a_1, a_2, a_3, b_1, b_2, b_3$ are distinct,
- $a_1, a_2, a_3, a'_1, a'_2, a'_3$ are distinct,
- $b_1, b_2, b_3, b'_1, b'_2, b'_3$ are distinct,
- G induces a triangle on each of $\{a_1, a_2, a_3\}$ and $\{b_1, b_2, b_3\}$, and
- for $i = 1, 2, 3$, a'_i is adjacent to a_i and b'_i is adjacent to b_i.

For each 15-tuple, do the following.

For $i = 1, 2, 3$ remove the subsets $N(a_i) \setminus \{a'_i, a_1, a_2, a_3\}$ and $N(b_i) \setminus \{b'_i, b_1, b_2, b_3\}$ from $V(G)$, and if $m_i = a'_i = b'_i$ then remove $N(m_i) \setminus \{a_i, b_i\}$ from $V(G)$. Now any smallest prism K that has $\{a_1, a_2, a_3, m_1, m_2, m_3, b_1, b_2, b_3\}$ as its frame and contains $\{a'_1, a'_2, a'_3, b'_1, b'_2, b'_3\}$ has a tidy frame.

Run the algorithm of 4.8 on G. This outputs $O(|V(G)|^6)$ subsets. For each subset Y output, do the following. Let G' be the graph induced by G on $V(G) \setminus Y$. Run the algorithm of 3.5, which takes $O(|V(G)|^5)$ time and outputs $O(|V(G)|^4)$ subsets. Let $X_1, \ldots, X_t$ be the subsets in the output. For every triple of subsets $(X_{i_1}, X_{i_2}, X_{i_3})$ from this list, do the following.

Let G_1 be the graph induced by G' on $V(G') \setminus X_{i_1}$. Find shortest paths S'_1 between a_1 and m_1 and T'_1 between m_1 and b_1 in G_1. Remove from $V(G_1)$ the set U_1 consisting of all vertices in $V(G_1) \setminus (S'^*_1 \cup T'^*_1 \cup \{a_1, b_1, m_1\})$ that have a neighbor in $S'^*_1 \cup T'^*_1$. Next, let G_2 be the graph induced by G_1 on $V(G_1) \setminus X_{i_2}$. Find shortest paths S'_2 between a_2 and m_2 and T'_2 between m_2 and b_2 in G_2. Remove from G_2 the set U_2 consisting of all vertices in $V(G_2) \setminus (S'^*_2 \cup T'^*_2 \cup \{a_2, b_2, m_2\})$ that have a neighbor in $S'^*_2 \cup T'^*_2$. Finally, let G_3 be the graph induced by G_2 on $V(G_2) \setminus X_{i_3}$. Find shortest paths S'_3 between a_3 and m_3 and between m_3 and b_3 in G_3. Finding these paths and removing these subsets takes quadratic time.

Finally, for $1 \leq i < j \leq 3$ test whether $V(S'_i) \cup V(T'_i) \cup V(S'_j) \cup V(T'_j)$ is the vertex set of an induced cycle in the graph G. If so, then output that the triangles $\{a_1, a_2, a_3\}$ and $\{b_1, b_2, b_3\}$ and the paths a_1-S'_1-m_1-T'_1-b_1, a_2-S'_2-m_2-T'_2-b_2 and a_3-S'_3-m_3-T'_3-b_3 form a prism and stop. Testing this takes quadratic time. So the time it takes to examine one 15-tuple is $O(|V(G)|^{20})$.

After examining all 15-tuples, output that G contains no smallest prism. The total running time is $O(|V(G)|^{35})$.

Now we need to prove that this algorithm is correct. Suppose there is a smallest prism K in G; let its half-paths be S_1, T_1, S_2, T_2, S_3 and T_3. For some 15-tuple chosen, $\{a_1, a_2, a_3, m_1, m_2, m_3, b_1, b_2, b_3\}$ is the frame of K, and $\{a'_1, a'_2, a'_3, b'_1, b'_2, b'_3\} \subseteq V(K)$. Therefore, after the first step of the algorithm runs for this 15-tuple, K has a tidy frame.

(1) *Let J be an almost clean smallest prism in a graph H and let S' be a good shortcut across a half-path S of J. Let $U \subset V(H)$ be the set of vertices not in $V(S')$ that have a neighbor in S'^*. Then the subgraph J' induced by H on $(V(J) \setminus V(S)) \cup V(S')$ is an almost clean smallest prism in the graph induced by H on $V(H) \setminus U$.*

Since S' is a good shortcut, the subgraph J' induced by H on $(V(J) \setminus V(S)) \cup V(S')$ is a prism. Since a good shortcut is a shortest path, $|V(J')| \leq$

$|V(J)|$, so J' is a smallest prism in H. We observe that since J' has the same frame as J, it has a tidy frame. Assume that there is a J'-major vertex $v \in V(H) \setminus U$. Then since $v \notin U$, $N(v) \cap S'^* = \emptyset$. Since v is J'-major, this means that $N(v) \cap (V(J') \setminus S'^*)$ is not a subset of the vertex set of a 3-vertex path in J'. Since $N(v) \cap (V(J') \setminus S'^*) = N(v) \cap (V(J) \setminus S^*)$, v is J-major, contradicting the fact that J is almost clean in H. Therefore, all J'-major vertices are contained in U. We observe also that U is disjoint from $V(J')$ because J' is a prism, so J' is an almost clean prism in the the graph induced by H on $V(H) \setminus U$. This proves (1).

By 4.8, we can choose Y such that G' contains no K-major vertices. Then K is almost clean in G', so by 3.5, we can choose X_{i_1} such that the graph G_1 contains K and there are no bad shortcuts across S_1 or T_1 in G_1. Let K_1 be the subgraph induced by G_1 on $V(K) \setminus (V(S_1) \cup V(T_1)) \cup V(S_1') \cup V(T_1')$. Since K_1 is not smaller than K in G, $S_1'^*$ is anticomplete to $T_1'^*$. It now follows from (1) applied twice (once for each half-path S_1 and T_1) that after removing the set U_1 from $V(G_1)$, K_1 is an almost clean smallest prism in G_1.

Therefore, by 3.5, we can choose X_{i_2} such that the graph G_2 contains K_1 and there are no bad shortcuts across S_2 or T_2 in G_2. Let K_2 be the subgraph induced by G_2 on $V(K_1) \setminus (V(S_2) \cup V(T_2)) \cup V(S_2') \cup V(T_2')$. Since K_2 is not smaller than K_1 in G, $S_2'^*$ is anticomplete to $T_2'^*$. Again it follows from (1) applied twice (once for each half-path S_2 and T_2) that after removing the set U_2 from $V(G_2)$, K_2 is an almost clean smallest prism in G_2.

Finally, by 3.5, we can choose X_{i_3} such that the graph G_3 contains K_2 and there are no bad shortcuts across S_3 or T_3 in G_3. Then it follows that the subgraph K_3 induced by G_3 on $V(K_2) \setminus (V(S_3) \cup V(T_3)) \cup V(S_3') \cup V(T_3')$ is a prism. Since K_3 is not smaller than K_2 in G, $S_3'^*$ is anticomplete to $T_3'^*$. We observe that since G_3 is an induced subgraph of G, a prism in G_3 is a prism in G, and that K_3 is the prism output by the algorithm when it considers $X_{i_1}, X_{i_2}, X_{i_3}$. This proves that the algorithm outputs a prism.

Conversely, if the algorithm outputs that there is a prism, then the properties of the 15-tuple and the conditions on the paths that the algorithm checks ensure that this output is actually a prism. □

Finally, we can put together the algorithms of 5.1 and 5.2 to construct the complete algorithm.

5.3. *There is an algorithm with the following specifications:*

- **Input:** *A graph G.*

- **Output:** *Either:*
  - a theta or a prism in G, or
  - a determination that there is no theta or prism in G.
- **Running Time:** $O(|V(G)|^{35})$.

Proof. First run the algorithm of 5.1. If it outputs a theta then output this theta and stop. Otherwise run the algorithm of 5.2. If it outputs a prism then output this prism and stop. Otherwise, output that there is no theta or prism in G.

The running time of this algorithm is the maximum of those of 5.1 and 5.2, which is $O(|V(G)|^{35})$. If this algorithm outputs either a theta or a prism, then it follows from 5.1 and 5.2 that it is correct. Conversely, if G contains a theta or a prism, then it contains a smallest theta or a smallest prism. Therefore, one of the algorithms of 5.1 or of 5.2 will output a theta or a prism, which will then be output by this algorithm. This proves 5.3. □

References

[1] M. Chudnovsky, G. Cornuéjols, X. Liu, P. Seymour, and K. Vušković, "Recognizing Berge graphs", *Combinatorica*, **25** (2005), no. 2, 143–186.

[2] M. Chudnovsky and P. Seymour, "The three-in-a-tree problem", *submitted for publication*.

[3] M. Conforti and M. R. Rao, "Testing balancedness and perfection of linear matrices", *Mathematical Programming* **61** (1993), 1–18.

[4] F. Maffray and N. Trotignon, "Algorithms for perfectly contractile graphs", *SIAM J. Discrete Math.* **19** (2005), no. 3, 553–574.