Detecting even holes

Maria Chudnovsky[1] Ken-ichi Kawarabayashi[2]
Paul Seymour[3]
Princeton University, Princeton NJ 08544

May 21, 2003; revised April 2, 2004

[1]This research was partially conducted during the period the author served as a Clay Mathematics Institute Research Fellow.
[2]Partially supported by the Japan Society for the Promotion of Science for Young Scientists
[3]Partially supported by ONR grant N00014-01-1-0608, and NSF grant DMS-0070912.

Abstract

Conforti, Cornuéjols, Kapoor and Vušković [4] gave a 73-page polynomial time algorithm to test whether a graph has an induced subgraph that is a cycle of even length. Here we provide another algorithm to solve the same problem. The differences are:

1. Our algorithm is simpler — we are able to search directly for even holes, while the algorithm of [4] made use of a structure theorem for even-hole-free graphs, proved in an earlier paper [3];

2. It is marginally faster — $O(n^{31})$ for an n-vertex graph (and we sketch another more complicated algorithm that runs in time $O(n^{15})$) while [4] appears to take about $O(n^{40})$; and

3. We can permit 0/1 weights on the edges and look for an induced cycle of even weight. Consequently we can test if a graph is "odd-signable".

1 Introduction

In [1, 2] two of us found an algorithm to test whether a graph contains an odd hole or antihole. There were a couple of novel techniques in that algorithm, and it seemed that they might be equally applicable when testing for even holes. It was already known that the problem of testing whether a graph contains an even hole is polynomially solvable [4], but that algorithm was very complicated, and we decided to try to adapt the method of [1, 2] to see if we could do it better. We did get an algorithm, not as nice as we hoped, but simpler than the algorithm of [4]. Incidentally, the complexity of finding the shortest even hole in a graph is still open as far as we know.

We begin with some definitions. All graphs in this paper are finite and simple. We denote the vertex-set of G by $V(G)$ and the edge-set by $E(G)$. A *path* in G means an induced subgraph that is a path (that is, it is non-null, connected, has no cycles, and has no vertex of degree > 2). Its *ends* are defined in the natural way, and for a path P with at least two vertices, its *interior* is the vertex set of the path obtained by deleting the ends of P. A *hole* in G means an induced subgraph that is a cycle. The *length* of a path or hole is the number of edges in it. (Note that cycles of length 3 count as holes. This is nonstandard, but convenient for us.) Two subsets $X, Y \subseteq V(G)$ are said to be *separate* if $X \cap Y = \emptyset$ and there do not exist $x \in X$ and $y \in Y$ such that x, y are adjacent.

We denote by Z_2 the two-element group with elements $\{0, 1\}$, using additive notation. A *signed graph* is a pair (G, γ), where G is a graph and γ is a function from $E(G)$ to Z_2. If $X \subseteq E(G)$, we define

$$\gamma(X) = \sum_{e \in X} \gamma(e).$$

If X is a path or hole in G, $\gamma(X)$ means $\gamma(E(X))$; and if $\gamma(X) = 1$, we say that X is γ-*odd* (and otherwise it is γ-*even*). We call $\gamma(X)$ the γ-*parity* of X. A graph G is said to be *odd-signable* if there exists γ such that there are no γ-even holes in the signed graph (G, γ). It is easy to show that the problem of testing whether a graph is odd-signable is polynomially equivalent to that of testing whether a signed graph (G, γ) has a γ-even hole.

In [4], Conforti, Cornuéjols, Kapoor and Vušković presented an algorithm to test whether G contains a hole of even length. It ran in time polynomial in the number of vertices of the input graph; they gave no explicit bound on its running time, but it seems to us to take time about $O(n^{40})$ for an n-vertex graph. Their algorithm relied on a decomposition theorem for graphs without even holes, proved in an earlier paper [3].

In this paper we give an algorithm which takes as input a signed graph (G, γ), and tests whether there is a γ-even hole in G. Its running time is $O(n^{31})$ (and we sketch some improvements that bring the running time down to $O(n^{15})$). It uses cleaning techniques to search directly for γ-even holes, and does not depend on any decomposition theorem. Consequently we can test if a graph is odd-signable in polynomial time.

Our algorithm was derived from the algorithm presented in [1, 2] to test whether a graph has an odd length hole or antihole (of length > 3). Searching for even holes is easier in some respects than searching for odd ones, and more difficult in some respects. For instance, it is still not known how to test in polynomial time whether a graph contains an odd hole of length > 3; the algorithm of [1, 2] tests for odd holes and antiholes simultaneously. On the other hand, there were some tricks that worked nicely for odd holes (testing first for pyramids, detecting clean shortest odd holes); they have analogues in the even hole case, but don't work so well.

The idea of the algorithm is as follows. Let (G, γ) be the input signed graph. Suppose that C is a γ-even hole in G of minimum length; how can we detect the presence of C? Let its length be t say. Let P be a path in G of length > 1 with ends u, v say in C, and let C_1, C_2 be the two paths of C between u, v, where C_1 is the shorter of the two. We say that P is a *shortcut* across C if its length is at most the length of C_1, and less than $t/4$. A shortcut P is *good* if its union with C_2 is another shortest γ-even hole; in other words, P has the same length and γ-parity as C_1, and its interior is separate from the interior of C_2. The shortcut is *bad* otherwise. A shortcut P is *shallow* if it is bad, the length of C_1 is at most one more than the length of P, and the union of P and C_2 is a hole. We first implement a "cleaning" subroutine; we generate polynomially many subsets of $V(G)$, such that it is guaranteed that one of them (say X) is disjoint from $V(C)$ and intersects enough of the bad shortcuts that, if any bad shortcuts still remain in $G \setminus X$, then the "worst" one of them is shallow. Now we examine all the graphs $G \setminus X$ (over all the choices of X that we generated, since we do not know which is the right one). If G has a γ-even hole, then in one of these graphs $G \setminus X$ there is either

- a shortest γ-even hole with no bad shortcut, or
- a shortest γ-even hole such that the worst shortcut over it is shallow.

In both kinds of graphs it is easy to detect that there is a γ-even hole. For example, for the first kind of graph we proceed as follows: we check all 8-tuples of vertices $v_1, \dots, v_8$; for each of the pairs $v_1v_2, v_2v_3, \dots, v_8v_1$ we find the shortest path joining the pair; and test whether the union of these eight paths is a γ-even hole. It can be shown that for graphs of the first kind, this algorithm will detect a γ-even hole. A similar method works for the second kind of graph.

We explain the cleaning subroutine in the next three sections, and then the two methods to detect γ-even holes in sections 5 and 6. In section 7 we sketch some variations which improve the running time to $O(|V(G)|^{15})$.

2 Major vertices

A *theta* in a graph G means an induced subgraph X of G with two nonadjacent vertices s, t and three paths P, Q, R, each between s, t, such that P, Q, R are internally disjoint, the union of every pair of them is a hole, and $X = P \cup Q \cup R$. A *prism* in G is an induced subgraph K in which there are three paths P_1, P_2, P_3, with the following properties:

- for $i = 1, 2, 3$, P_i has length > 0; let its ends be a_i, b_i
- P_1, P_2, P_3 are pairwise disjoint, and $V(K) = V(P_1 \cup P_2 \cup P_3)$
- for $1 \leq i < j \leq 3$, there are precisely two edges between $V(P_i)$ and $V(P_j)$, namely a_ia_j and b_ib_j.

It is easy to see that in any theta and in any prism, at least one of the holes has even γ-parity, for any choice of γ.

Throughout this section, let (G, γ) be a signed graph with a γ-even hole, and let C be such a hole with minimum length, length t say. We call C a *shortest γ-even hole.* We observe:

2.1 *Let $C_1, \dots, C_k$ be a list of holes in G with k odd, such that every edge is in an even number of $C_1, \dots, C_k$. Then one of them is γ-even, and therefore has length at least t. (In particular, for every theta or prism in G, at least one of its holes has length $\geq t$.) Also, let $C_1, \dots, C_k$ be a list of holes with k odd such that every edge is in an even number of $C, C_1, \dots, C_k$. Then one of $C_1, \dots, C_k$ is γ-even, and therefore has length at least t.*

Proof. Suppose there is a list as in the first assertion. Since every edge is in an even number of $C_1, \dots, C_k$, the sum of their parities is 0, and so one of them is γ-even, and therefore has length at least t. The proof of the second assertion is similar, using that C is γ-even. This proves 2.1. ∎

2.2 *Let C' be a hole in G of length $\leq t$, and let $v \in V(G) \setminus V(C')$. Then either:*

- *there is an edge xy of C' such that v has no neighbours in $V(C') \setminus \{x, y\}$, or*
- *C' has length t and there are three consecutive vertices x, y, z of C' such that v has no neighbours in $V(C') \setminus \{x, y, z\}$, or*
- *v has an even number of neighbours in $V(C')$ if and only if C' is γ-even.*

Proof. Let us say a path of C' with both ends adjacent to v and no internal vertex adjacent to v is a *gap*. We may assume that the first outcome of the theorem does not hold. Consequently, every edge of C' belongs to a unique gap. Every gap can be completed to a hole, via a two-edge path with middle vertex v. The sum of the γ-parities of these holes equals the γ-parity of C', and the number of these holes is the number of neighbours of v in C'; so if all these holes are γ-odd then the third outcome of the theorem holds. We may therefore assume that one of them is γ-even. Consequently it has length $\geq t$, and so the corresponding gap has length $\geq t - 2$. It follows that C' has length t and the second outcome holds. This proves 2.2. ∎

Let us say a vertex $v \in V(G) \setminus V(C)$ is *major* if there are three of its neighbours in C that are pairwise nonadjacent. Our first task is to clean away major vertices, and in fact this is the most difficult step of the entire algorithm. Let us assign an orientation to C called *clockwise*, and let C have vertices $c_1, \dots, c_t$ in clockwise order. If $a, b \in V(C)$ are distinct, we denote by $C(a, b)$ the subgraph of C consisting of all the vertices and edges of C traversed as we move from a to b along C in the clockwise direction. So if a, b are nonadjacent then $C(a, b)$ and $C(b, a)$ are the two paths of C between a and b.

Let u, v be nonadjacent major vertices. A *gate* for the ordered pair (u, v) is an edge xy of C with the following properties:

- y follows x as C is traversed in clockwise direction
- uy, vx and at least one of ux, vy are edges
- there is a vertex $z \in V(C) \setminus \{x, y\}$ such that u has no neighbours in the interior of $C(y, z)$, and v has no neighbours in the interior of $C(z, x)$.

A vertex z as above is called a *divider* for the gate.

2.3 *For every pair u, v of distinct nonadjacent major vertices, either there is a unique gate for (u, v) and none for (v, u), or vice versa.*

Proof.

(1) *There is at most one gate for* (u, v)*.*

For suppose that there are two, say c_1c_2 and c_ic_{i+1} where $2 \leq i < t$. Since u is major, it has a neighbour different from c_1, c_2, c_i, c_{i+1}, say c_j. Similarly v has a neighbour c_h in C different from c_1, c_2, c_i, c_{i+1}. Since c_1c_2 is a gate for (u, v), and uc_{i+1}, vc_h are edges, it follows that $h \leq i + 1$ and consequently $h < i$; and since uc_j, vc_i are edges it follows similarly that $j > i + 1$. But since uc_j, vc_h are edges, this contradicts that c_ic_{i+1} is a gate for (u, v). This proves (1).

(2) *There is not both a gate for* (u, v) *and a gate for* (v, u)*.*

For suppose that c_1c_2 is a gate for (u, v), and some c_ic_{i+1} is a gate for (v, u) where $1 \leq i < t$. Let c_k be a divider for the first gate. Since v has no neighbours in the interior of $C(c_k, c_1)$ and c_{i+1} is a neighbour of v and $i < t$, it follows that $k \geq i + 1$; and since u has no neighbours in the interior of $C(c_2, c_k)$ and c_i is a neighbour of u, it follows that $i \leq 2$. Since u is major, it follows from 2.2 that there is a neighbour c_h of u with $4 \leq h \leq t$; and similarly there is a neighbour c_j of v with $4 \leq j \leq t$. Since c_1c_2 is a gate for (u, v), it follows that $h \geq j$; and since c_ic_{i+1} is a gate for (v, u) and $i = 1$ or 2, it follows that $h \leq j$. Consequently $h = j$, and so both c_h and c_j are unique. Since u, v both have at least four neighbours in C by 2.2, it follows that they both have exactly four neighbours in C, namely c_1, c_2, c_3 and c_h. Since u is major it follows that $4 < h < t$ and in particular, $t \geq 6$; but then the vertices u, v, c_1, c_3, c_h induce a theta in which all holes have length 4. Since $4 < n$, i this constradicts 2.1. This proves (2).

It remains to show that there is a gate for at least one of $(u, v), (v, u)$. Let us say a *gap* is a minimal path in C containing a neighbour of u and a neighbour of v. It follows that every edge of C is in at most one gap, and consequently any two distinct gaps share no edges, and any common vertex is an end of both. Every gap P is the interior of a path between u, v, say P^+.

(3) *If* P_1, P_2 *are distinct gaps and* P_1^+, P_2^+ *have the same* γ*-parity, then* P_1, P_2 *are disjoint and an end of* P_1 *is adjacent to an end of* P_2*.*

For suppose first that they are not disjoint, and therefore share an end, c_i say. We may assume that P_1 is $C(c_h, c_i)$ and P_2 is $C(c_i, c_j)$, where $1 \leq h \leq i \leq j \leq t$. Since c_i belongs to two gaps, it is not adjacent to both u and v and so $h < i < j$. We may assume that u is adjacent to c_h, c_j and v to c_i. But then $C(c_h, c_j)$ can be completed to a hole via c_j-u-c_h, and it is γ-even since its γ-parity is the sum of the parities of P_1^+ and P_2^+; and its length is $< t$ since u is major, a contradiction. So P_1, P_2 are disjoint. Suppose there is no edge between them. Then $P_1^+ \cup P_2^+$ is a γ-even hole, and therefore has length $\geq t$. Since P_1, P_2 are disjoint and there are no edges between them, and C has length t, it follows that there are two vertices a, b of C, nonadjacent, so that P_1 is the interior of $C(a, b)$ and P_2 the interior of $C(b, a)$. Since u has exactly two neighbours in $P_1 \cup P_2$, and has at least four in C, it follows that ua, ub are edges, and similarly so are va, vb. The subgraph induced on $\{u, v\} \cup V(C(a, b))$ is a prism, and all its holes are strictly shorter than C, contrary to 2.1. This proves (3).

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

For let the gaps be $P_1, \dots, P_k$, numbered in their order on C, and suppose that $k \geq 7$. It follows that P_1, P_5 are disjoint and there are no edges between them, and the same holds for P_2, P_5. By (3), P_1^+, P_5^+ have opposite γ-parity, and so do P_2^+, P_5^+; and consequently P_1^+, P_2^+ have the same γ-parity. Similarly P_2^+, P_3^+ have the same γ-parity, and so do all such consecutive pairs, and therefore so do P_1^+, P_5^+, a contradiction. This proves (4).

(5) *There are at most five gaps.*

For assume there are exactly six, say $P_1, \dots, P_6$ in order. By (3) we deduce that P_1^+, P_4^+ have opposite γ-parity, and so do P_2^+, P_5^+ and P_3^+, P_6^+. Since $P_1^+, \dots, P_6^+$ do not all have the same γ-parity, there are two consecutive with opposite γ-parity, say P_1^+, P_2^+. It follows that P_2^+, P_4^+ have the same γ-parity, and so by (3) P_2, P_4 are disjoint and there is an edge between them; and so P_3 is not disjoint from both of P_2, P_4, and therefore by (3), P_3^+ has opposite γ-parity from P_2^+. We have shown then that if P_1^+, P_2^+ have opposite γ-parity then so have P_2^+, P_3^+. Consequently every consecutive pair of gaps are the interiors of $u - v$ paths of opposite γ-parity. We also showed that there is an edge between P_2, P_4, and therefore P_3 shares a vertex with one of P_2, P_4. Consequently P_3 has length > 0, and so it has length 1 and shares an end with P_2 and shares the other end with P_4. The same holds for every other gap, and so $t = 6$. But then u has exactly three neighbours in C, contrary to 2.2. This proves (5).

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

For assume there are exactly five, say $P_1, \dots, P_5$ in order. Since there is an odd number of gaps, it follows that u, v have a common neighbour in C, and so some gap has length 0, say P_5. So P_2, P_5 are disjoint and there is no edge between them, and P_2^+, P_5^+ therefore have opposite γ-parity, by (3); and similarly P_3^+, P_5^+ have opposite γ-parity. Also P_1, P_4 are disjoint and there is no edge between them, so P_1^+, P_4^+ have opposite γ-parity; and therefore from the symmetry between P_1 and P_4, we may therefore assume that P_1^+, P_5^+ have opposite γ-parity. It follows that P_1^+, P_2^+, P_3^+ all have the same γ-parity, so by (3) they are pairwise disjoint and every pair of them is joined by an edge, which is impossible. This proves (6).

(7) *There are at most three gaps.*

For assume there are exactly four, say $P_1, \dots, P_4$ in order. Suppose first that at least three of $P_1^+, \dots, P_4^+$ have the same γ-parity. Then these three are pairwise disjoint and pairwise joined by edges, by (3), and so have union C; but since u has only one neighbour in each of them, it has only three neighbours in C, contrary to 2.2. So two of $P_1^+, \dots, P_4^+$ have γ-parity 0 and the other two have γ-parity 1. Suppose that P_1^+, P_3^+ have equal γ-parity, and hence P_2^+, P_4^+ have the other γ-parity. By (3) there is an edge between P_1 and P_3, and also one between P_2, P_4. Consequently we may assume that $P_1, \dots, P_4$ all have a vertex in $\{c_1, c_2, c_3\}$, and so no subpath of $C(c_4, c_t)$ is a gap. It follows that not both u, v have neighbours in $C(c_4, c_t)$, a contradiction since they are both major. So we may assume that P_1^+, P_2^+ have the same γ-parity, and therefore P_3^+, P_4^+ have the other γ-parity.

Consequently P_1, P_2 are disjoint and are joined by an edge and so are P_3, P_4. These two edges are necessarily disjoint; so we may assume that they are c_1c_2 and c_ic_{i+1} for some i with $3 \le i < t$. Thus c_1 is an end of P_1, and P_1 is contained in $C(c_{i+1}, c_1)$; c_2 is an end of P_2 and P_2 is contained in $C(c_2, c_i)$; and similarly for P_3, P_4. Since the interior of $C(c_2, c_i)$ includes no gap, not both u, v have neighbours in it, and similarly they do not both have neighbours in the interior of $C(c_{i+1}, c_1)$. But since they are both major, they both have neighbours in C different from c_1, c_2, c_i, c_{i+1}; so we may assume that u has neighbours in the interior of $C(c_2, c_i)$, and not in the interior of $C(c_{i+1}, c_1)$, and vice versa for v. Since P_2 is a gap and therefore contains a neighbour of v, it follows that c_2 is adjacent to v, and similarly vc_i, uc_1, uc_{i+1} are edges. Hence $C(c_2, c_i)$ can be completed to a hole via c_i-v-c_2, and this hole (say C') has length $< t$. But u has the same neighbours in C and in C', except for c_1, c_{i+1}, and therefore it has an even number in C', since it has an even number in C by 2.2. By 2.2 applied to C', we deduce u has exactly two neighbours in C' (and thus in $C(c_2, c_i)$) and they are adjacent. Similarly v has exactly two neighbours in $C(c_{i+1}, c_1)$ and they are adjacent. Since c_1-c_2 is not a gap, it follows that one of u, v is adjacent to both c_1, c_2, and similarly for c_i, c_{i+1}. We may therefore assume that u is adjacent to c_2 and so its only neighbours in C are c_1, c_2, c_3, c_{i+1}. Since u is major, it follows that $i > 3$, and so u is nonadjacent to c_i. Hence v is adjacent to c_i and c_{i+1}. The subgraph induced on $\{u, v, c_2, c_3, \dots, c_{i+1}\}$ is a prism and all its holes have length $< t$, contrary to 2.1. This proves (7).

(8) *There are exactly three gaps.*

For assume there are at most two. Then there exist two vertices of C, say c_1, c_i, such that every gap contains at least one of them. Consequently not both u, v have neighbours in the interior of $C(c_1, c_i)$, and similarly for $C(c_i, c_1)$. Since they do both have more neighbours, we may assume that all neighbours of u in C belong to $C(c_1, c_i)$, and all neighbours of v belong to $C(c_i, c_1)$. We may also assume that vc_1, uc_i are edges. Choose h, j with $1 \le h, j \le t$ minimum such that uc_h and vc_j are edges; then $1 \le h \le i \le j \le t$. Since all neighbours of u in C belong to $C(c_h, c_i)$, and it has at least four such neighbours, it follows that $i \ge h+3$, and similarly $j \le t-2$. But then every edge of G is in an even number of the holes C, v-c_1-c_2-$\cdots$-c_j-v, u-c_i-c_{i+1}-$\cdots$-c_t-c_1-$\cdots$-c_h-u, v-c_1-$\cdots$-c_h-u-c_i-$\cdots$-c_j-v, and the latter three all have length $< t$, contrary to 2.1. This proves (8).

From (8), there are two gaps with the same γ-parity, and so by (3) these gaps are disjoint and joined by an edge. We may assume the edge is c_tc_1 and there exist h, j with $1 \le h < j \le t$ such that $P_1 = C(c_j, c_t)$ and $P_2 = C(c_1, c_h)$. In particular, c_1, c_t are both adjacent to one of u, v. Since u, v both have neighbours in $C(c_2, c_{t-1})$ and hence P_3 is included in this path, it follows that the path c_t-c_1 is not a gap, and so one of u, v, say u, is adjacent to both of c_t, c_1. Consequently v is adjacent to c_h, c_j. Now $C(c_{h+1}, c_{j-1})$ includes at most one gap, and u, v both have at least two neighbours in it, so there exists i with $h+2 \le i \le j-2$ such that one of u, v has no neighbours in $C(c_{h+1}, c_{i-1})$ and the other has no neighbours in $C(c_{i+1}, c_{j-1})$. We may assume (by reversing the orientation of C if necessary) that u has no neighbours in $C(c_{h+1}, c_{i-1})$ and v has none in $C(c_{i+1}, c_{j-1})$. By reducing i if necessary, we may also assume that v is adjacent to c_i. Since u has at least two neighbours in $C(c_i, c_j)$, and v is adjacent to both c_i, c_j, it follows that $C(c_i, c_j)$ includes at least two gaps, neither of which is P_2. Consequently one of them is P_1, and so $j = t$. But then c_tc_1 is a gate for (u, v), and the theorem holds. This proves 2.3. ∎

2.4 *Let u, v, w be distinct major vertices, such that w is nonadjacent to both u and v. Then there is a vertex in C adjacent to all of u, v, w.*

Proof. We assume for a contradiction that there is no such vertex.

(1) *No edge of C is a gate for one of $(u, w), (w, u)$ and a gate for one of $(v, w), (w, v)$.*

For suppose that $c_t c_1$ is such an edge. Consequently one of c_t, c_1 is adjacent to both u and w, and one is adjacent to both v, w. Since neither of c_t, c_1 is adjacent to all three of u, v, w, we may assume that c_1 is adjacent to u, w and not to v, and c_t is adjacent to v, w and not to u. Consequently $c_t c_1$ is not a gate for (w, u), so it is a gate for (u, w), and similarly it is a gate for (w, v). Choose i, j with $2 \le i \le j \le t - 1$, such that wc_i, wc_j are edges, with i minimum and j maximum. Since w is major it follows that $j \ge i + 2$. Since $c_t c_1$ is a gate for (u, w) it follows that u has no neighbours in $C(c_2, c_{j-1})$, and similarly v has none in $C(c_{i+1}, c_{t-1})$. Thus u, v have no common neighbour in C, and hence u, v are adjacent, by 2.3. Choose h with $1 \le h \le i$ maximum such that vc_h is an edge, and k with $j \le k \le t$ minimum such that uc_k is an edge. Since v is major it follows that $h \ge 4$, and similarly $k \le t - 3$. But then the three paths u-c_1-w, u-v-c_h-$\cdots$-c_i-w and u-c_k-$\cdots$-c_j-w form a theta in which all three holes have length $< t$, contrary to 2.1. This proves (1).

(2) *There do not exist two consecutive edges of C such that one is a gate for either (u, w) or (w, u), and the other is a gate for either (v, w) or (w, v).*

For assume that, say, $c_t c_1$ is a gate for one of $(u, w), (w, u)$, and $c_1 c_2$ is a gate for one of $(v, w), (w, v)$. Suppose first that u is nonadjacent to c_1. Hence $c_t c_1$ is not a gate for (u, w), and so it is a gate for (w, u). It follows that wc_t is an edge, and w is nonadjacent to c_2 (because u has some neighbour c_i where $i \ge 3$, and therefore w has no neighbours in the interior of $C(c_1, c_i)$). The first implies that $c_1 c_2$ is not a gate for (v, w), and the second implies that $c_1 c_2$ is not a gate for (w, v), a contradiction. It follows that uc_1 is an edge, and similarly vc_1 is an edge. Since c_1 is not adjacent to all of u, v, w it follows that w is nonadjacent to c_1. So $c_t c_1$ is not a gate for (w, u), and therefore it is a gate for (u, w), and similarly $c_1 c_2$ is a gate for (w, v). So wc_t, wc_2, uc_t, vc_2 are edges. Hence c_t is nonadjacent to v since it is not adjacent to all of u, v, w, and similarly c_2 is nonadjacent to u. Since w is major, it has a neighbour c_i in $C(c_3, c_{t-1})$. Since $c_t c_1$ is a gate for (u, w) it follows that u has no neighbours in $C(c_2, c_{i-1})$, and similarly v has none in $C(c_{i+1}, c_t)$. There is therefore no gate for (u, v) or (v, u), and so they are adjacent by 2.3. But then there are five holes in the subgraph induced on $\{u, v, w, c_t, c_1, c_2\}$, all of length $< t$, and every edge is in an even number of them, contrary to 2.1. This proves (2).

In view of 2.3 and (2), we may assume that $c_t c_1$ is a gate for (u, w), and some $c_i c_{i+1}$ is a gate for one of $(v, w), (w, v)$, where $2 \le i \le t - 2$. In particular, uc_1, wc_t are edges.

(3) *$c_i c_{i+1}$ is a gate for (w, v). Consequently, vc_i, wc_{i+1} are edges; u has no neighbours in $C(c_2, c_i)$; and v has no neighbours in $C(c_1, c_{i-1})$.*

For suppose $c_i c_{i+1}$ is not a gate for (w, v); then $c_i c_{i+1}$ is a gate for (v, w). In particular, vc_{i+1} and wc_i are edges. Since w is adjacent to c_t and $c_i c_{i+1}$ is a gate for (v, w), it follows that v has no neighbours in $C(c_{i+2}, c_{t-1})$. We claim also that v is not adjacent to c_t. For suppose it is; then u is not adjacent to c_t, since c_t is not adjacent to all of u, v, w, and w is not adjacent to c_1 since c_ic_{i+1} is a gate for (v, w); and this contradicts that c_tc_1 is a gate for (u, w). So v is not adjacent to c_t. Choose h with $1 \leq h \leq t$ minimum such that v is adjacent to c_h; then $1 \leq h \leq i$, and all neighbours of v in C belong to $C(c_h, c_{i+1})$ (and consequently $h \leq i-3$). Similarly there is a neighbour c_j of u with $i+1 \leq j \leq t-3$ such that all neighbours of u in C belong to $C(c_j, c_1)$. Thus there is no gate for (u, v) or for (v, u), and therefore u, v are adjacent by 2.3. Since c_tc_1 is a gate for (u, w), it follows that w has no neighbours in $C(c_{j+1}, c_{t-1})$, and similarly it has none in $C(c_{h+1}, c_{i-1})$. Since w is major, it has a neighbour in C different from c_t, c_1, c_i, c_{i+1}, and from the symmetry between c_tc_1 and c_ic_{i+1}, we may assume that wc_g is an edge, where $2 \leq g \leq h$. Choose g maximum with this property. Choose k with $1 \leq k \leq t$, maximum such that uc_k is an edge; then $k \geq j+2 \geq i+3$ since u has at least four neighbours in C. Let P_1 be the path w-c_g-$\cdots$-c_h-v, and P_2 a path between w, v with interior in $\{c_i, c_{i+1}\}$, and P_3 the path w-c_t-c_{t-1}-$\cdots$-c_k-u-v; then the union of these three paths is a theta, and the union of any two of them is a hole of length $< t$, contrary to 2.1. This proves the first assertion of (3). It follows that vc_i, wc_{i+1} are edges. Since w is adjacent to c_{i+1} and c_tc_1 is a gate for (u, w), it follows that u has no neighbours in $C(c_2, c_i)$; and since w is adjacent to c_t and c_ic_{i+1} is a gate for (w, v), it follows that v has no neighbours in $C(c_1, c_{i-1})$. This proves (3).

(4) *w has no neighbours in $C(c_{i+2}, c_{t-1})$.*

For suppose it does, and choose k' with $i+2 \leq k' \leq t-1$, such that w is adjacent to $c_{k'}$. Since $c_{k'}$ is not adjacent to all of u, v, w, we may assume from the symmetry that $c_{k'}$ is not adjacent to u. Choose k minimum such that $i+2 \leq k \leq t-1$ and w is adjacent to c_k; then it follows that c_k is not adjacent to u (because c_tc_1 is a gate for (u, w)). Since c_tc_1 is a gate for (u, w), it follows that u has no neighbours in $C(c_2, c_{k-1})$; and consequently $k \leq t-2$. Since c_ic_{i+1} is a gate for (w, v), it follows that v has no neighbours in $C(c_{k+1}, c_{i-1})$. Since c_k is not adjacent to all of u, v, w, it follows that u, v have no common neighbour in C, and so they are adjacent, by 2.3. Choose j with $1 \leq j \leq t$ maximum such that v is adjacent to c_j; then $i+3 \leq j \leq k$. Let P_1 be the path v-c_j-$\cdots$-c_k-w, let P_2 be a path from v to w with interior in $\{u, c_t, c_1\}$, and let P_3 be the path from v to w with interior in $\{c_i, c_{i+1}\}$. The union of these three paths is a theta, and all the holes in it have length $< t$, contrary to 2.1. This proves (4).

(5) *w has exactly two neighbours in $C(c_1, c_i)$, and they are adjacent.*

For u, v both have neighbours in $C(c_{i+2}, c_{t-1})$, and so there is a path P between u, v with interior in $\{c_{i+2}, \ldots, c_{t-1}\}$ (if u, v are adjacent then P has length 1.) Let C' be the hole formed by the union of c_1-u-P-v-c_i and $C(c_1, c_i)$; its length is $< t$, so it has odd γ-parity. But by 2.2, w has an even number of neighbours in C, and it has exactly the same neighbours in C' except for c_{i+1}, c_t, and therefore it has an even number in C'. By 2.2, it has exactly two adjacent in C'. This proves (5).

(6) *u, v are nonadjacent.*

Let the neighbours of w in $C(c_1, c_i)$ be c_h, c_{h+1}. Since w is major, it has neighbours in C different from c_t, c_1, c_i, c_{i+1}, so we may assume that $h > 1$. So w is not adjacent to c_1. Since c_tc_1 is a gate for (u, w), it follows that u is adjacent to c_t. If u is adjacent to v, then the subgraph induced on

$\{u, v, w, c_t, c_1, \dots, c_i\}$ is a prism, and all its holes have length $< t$, a contradiction. This proves (6).

This restores the symmetry between u, v, w. By 2.3, (1) and (2) there is a gate for one of $(u, v), (v, u)$, different from and disjoint from both $c_t c_1$ and $c_i c_{i+1}$; and since u, v have a common neighbour in this gate, it therefore belongs to $C(c_{i+1}, c_t)$. Let $c_k c_{k+1}$ be this gate, where $i + 2 \leq k \leq t - 2$. Since v has a neighbour c_i in the interior of $C(c_1, c_k)$, it follows that $c_k c_{k+1}$ is not a gate for (u, v), and therefore it is a gate for (v, u). So all neighbours of v in C belong to $C(c_i, c_{k+1})$, and all neighbours of u in C belong to $C(c_k, c_1)$; and uc_k is an edge. Let the neighbours of w in $C(c_1, c_i)$ be c_h, c_{h+1}. As before, we may assume from the symmetry that $h > 1$, and therefore u is adjacent to c_t. Choose j with $1 \leq j \leq k$ maximum such that v is adjacent to c_j; so $j \geq i + 2$. Then the three paths $c_1\text{-}\cdots\text{-}c_h$, $c_t\text{-}w$, and $u\text{-}c_k\text{-}\cdots\text{-}c_j\text{-}v\text{-}c_i\text{-}\cdots\text{-}c_{h+1}$ form a prism, and all its holes have length $< t$, a contradiction. This proves 2.4. ∎

We use 2.4 for the following algorithm.

2.5 *There is an algorithm with the following specifications:*

- **Input** *A signed graph* (G, γ).
- **Output** *A sequence of subsets* $X_1, \dots, X_r$ *of* $V(G)$*, with* $r \leq |V(G)|^9$*, such that for every shortest* γ*-even hole* C *in* G*, one of* $X_1, \dots, X_r$ *is disjoint from* $V(C)$ *and includes all major vertices for* C*.*
- **Running time** $O(|V(G)|^{10})$.

Proof. The algorithm is as follows. If G has at most two vertices, we output the null sequence; this output evidently has the desired properties, and so henceforth we assume that G has at least three vertices. In the first phase, we enumerate all 9-tuples $v_1, \dots, v_9$ of distinct vertices of G. For each such 9-tuple, let X be the set of all vertices of G that are different from $v_3, \dots, v_9$ and are either adjacent to one of v_4, v_5, v_8 or are adjacent to both of v_1, v_2. Let $X_1, \dots, X_a$ be the different subsets X generated in this phase.

In the second phase, we generate all 8-tuples $v_1, \dots, v_8$ of vertices (not necessarily distinct). For each such 8-tuple, let Y be the set of all vertices adjacent to all of v_1, v_2, v_3; if $v_8 \notin Y$ let Z be the set of all vertices in $V(G) \setminus (Y \cup \{v_1, \dots, v_8\})$ that are adjacent to all of $Y \cup \{v_8\}$, and if $v_8 \in Y$ let $Z = \emptyset$; and let $X = Y \cup Z \cup \{v_8\}$. Let $X_{a+1}, \dots, X_b$ be the subsets X generated in this phase.

We output the list $X_1, \dots, X_b, \emptyset$. That concludes the description of the algorithm; now we prove that it works correctly. For an n-vertex graph, the total number of subsets that we output is

$$n(n-1)(n-2)\cdots(n-8) + n^8 + 1 \leq n^9,$$

as claimed; and the running time is indeed $O(n^{10})$, because it takes linear time to process each 9-tuple and quadratic time for each 8-tuple. (If the running time was critical, we could arrange it more efficiently, but it turns out that improvements here make only a negligible improvement in the running time of the complete algorithm.)

It remains to check that the sequence $X_1, \dots, X_b, \emptyset$ that we generate has the property claimed for the output. Let C be a shortest γ-even hole in G. If there are no major vertices for C then the sequence works because it contains the empty set. Suppose next that there are two nonadjacent major vertices, say v_1, v_2. By 2.3, there are seven distinct vertices $v_3, \ldots, v_9$ of C, such that v_3-v_4-v_5-v_6 and v_7-v_8-v_9 are paths, and v_1, v_2 have no common neighbours in C except possible v_4, v_5, v_8. Let X be the corresponding subset generated in phase 1 of the algorithm when we come to examine this 9-tuple. Since every vertex in X is either adjacent to one of v_4, v_5, v_8 or is adjacent to both of v_1, v_2, and $v_3, \ldots, v_9 \notin X$, it follows that X contains no vertex of C. On the other hand, we claim that every major vertex belongs to X. Certainly v_1, v_2 each are adjacent to one of v_4, v_5, v_8; indeed, by 2.3 they have a common neighbour in $V(C)$, which must be one of v_4, v_5, v_8. Consequently $v_1, v_2 \in X$. Let $v_3 \neq v_1, v_2$ be a major vertex. If v_3 is adjacent to both v_1, v_2 then $v_3 \in X$ from the definition of X; and otherwise v_1, v_2, v_3 have a common neighbour in $V(C)$ by 2.4, which therefore must be one of v_4, v_5, v_8, and again $v_3 \in X$. This proves that all major vertices belong to X, so in this case the claim holds. Finally we assume there is a major vertex, but all major vertices are pairwise adjacent. Since there is a major vertex, it follows that C has length ≥ 6. Choose v_1, v_2, v_3 in C, pairwise nonadjacent and with as many common neighbours as possible. Let Y be the set of their common neighbours. Since there is a major vertex, it follows that $Y \neq \emptyset$. Let v_8 be a major vertex, chosen nonadjacent to one of v_1, v_2, v_3 if possible. If $v_8 \in Y$ then Y contains all major vertices and is disjoint from C, and Y belongs to the output list, so the claim holds. Finally we assume that $v_8 \notin Y$. From the choice of v_1, v_2, v_3 it follows that there do not exist three nonadjacent vertices in C all adjacent to all of $Y \cup \{v_8\}$. Consequently there are four vertices v_4, v_5, v_6, v_7 in C such that every vertex in C that is adjacent to all of $Y \cup \{v_8\}$ is one of these four. Let Z, X be the sets defined in phase 2 of the algorithm when we come to examine the 8-tuple $v_1, \ldots, v_8$. Certainly all members of Y are major, and therefore not in C. Every vertex in Z is adjacent to all of $Y \cup \{v_8\}$, and different from $v_1, \ldots, v_7$, and is therefore not in C; so X is disjoint from C. It remains to show that X contains all major vertices. Let w be a major vertex. If $w \in Y \cup \{v_8\}$ then $w \in X$ as required, so we assume not; but w is adjacent to all of $Y \cup \{v_8\}$, since every two major vertices are adjacent, and therefore $w \in Z \subseteq X$, as required. This proves 2.5. ∎

3 Shortcuts

For any two vertices u, v of any graph H, we denote by $d_H(u, v)$ the length of the shortest path of H between u, v. Let (G, γ) be a signed graph, and let C be a shortest γ-even hole in G, of length t. A *shortcut across* C is a path P in G with distinct nonadjacent ends $u, v \in V(C)$, such that P has length $< t/4$, and length $\leq d_C(u, v)$. A shortcut P is *good* if its union with one of $C(u, v), C(v, u)$ is another γ-even hole of length t, and *bad* otherwise. In other words, assume that $C(u, v)$ has length at most that of $C(v, u)$; then P is good if all the following are satisfied:

- P and $C(u, v)$ have the same length,
- no internal vertex of P belongs to $C(v, u)$,
- no internal vertex of P has a neighbour in the interior of $C(v, u)$, and
- P and $C(u, v)$ have the same γ-parity.

If P is a shortcut, we say P is *shallow* if it is bad, the interior of P is separate from the interior of $C(v, u)$, and the length of P is at least $d_C(u, v) - 1$. We say it is *deep* if it is bad and there is no shallow shortcut with the same interior as P. We say a shortcut P with vertices u-p_1-$\cdots$-p_k-v is *clear* if it is deep (and therefore bad), and no internal vertex of P belongs to C, and there are two subpaths Q_1, Q_2 of C with the following properties:

- p_1 is adjacent to every vertex of Q_1, and p_k is adjacent to every vertex of Q_2, and there are no other edges between $\{p_1, \dots, p_k\}$ and $V(C)$
- Q_1, Q_2 are disjoint; they both have length ≤ 2; and they both have even length or both have length 1
- if $k = 1$ then Q_1, Q_2 both have length at most 1.

We call Q_1, Q_2 the *attachments paths* of P. A clear shortcut P is said to be of *even* (respectively, *odd*) *type* depending whether its attachment paths Q_1, Q_2 have even or odd length.

Let P, P' be bad shortcuts across C, with ends u, v and u', v' respectively. We say that P' is *worse* than P if either

- the length of P' is strictly less than that of P, or
- P' and P have the same length, and $d_C(u, v) < d_C(u', v')$.

If P is a bad shortcut and there is no bad shortcut that is worse than P, we say P is a *worst* shortcut. In this section we analyze the possible types of a worst shortcut.

3.1 *Let (G, γ) be a signed graph, and let C be a shortest γ-even hole in G, and let P be a worst shortcut across C. Then either:*

- *P has exactly one internal vertex, and it is major, or*
- *P is clear, or*
- *P is shallow.*

Proof. For a contradiction, we suppose that P does not satisfy the theorem. Let P have vertices u-p_1-$\cdots$-p_k-v, and let C have vertices c_1-$\cdots$-c_t, in order. We may assume that $C(u, v)$ has length at most that of $C(v, u)$.

(1) *$p_1, \dots, p_k \notin V(C)$.*

For assume that some $p_i \in V(C)$. Since u-p_1-$\cdots$-p_i is not a bad shortcut, it follows that $d_C(u, p_i) \leq i$, and similarly $d_C(p_i, v) \leq k - i + 1$. Consequently $d_C(u, v) \leq k + 1$, and since P has length $k + 1$, we have equality throughout. In particular, $C(u, v)$ has length $k + 1$, and p_i belongs to a minimum length path of C between u, v, and so we may assume that $p_i \in V(C(u, v))$. Since u-p_1-$\cdots$-p_i is not a bad shortcut, and it has the same length as $C(u, p_i)$, it follows that it also has the same γ-parity as $C(u, p_i)$, and there are no edges between its interior and the interior of $C(p_i, u)$. Similar statements hold for the path p_i-$\cdots$-p_k-v; but this contradicts that P is a bad shortcut. This proves (1).

(2) *$k \geq 2$.*

For suppose that $k = 1$. Then p_1 is adjacent to both u, v. If it has exactly two neighbours in

C, then P is either shallow or clear, a contradiction. Assume that it has exactly three; then by 2.2 they are consecutive, say c_1, c_2, c_3, where C has vertices $c_1, \dots, c_t$ in order. Since u, v are nonadjacent, it follows that $\{u, v\} = \{c_1, c_3\}$; and since P is a bad shortcut, and has the same length as $C(c_1, c_3)$ and p_1 has no neighbours in the interior of $C(c_3, c_1)$, it follows that P and $C(c_1, c_3)$ have opposite γ-parity. Consequently one of the two holes p_1-c_1-c_2-p_1, p_1-c_2-c_3-p_1 has even γ-parity, and therefore $t = 3$, contradicting that u, v are nonadjacent. So p_1 has at least four neighbours in C. By 2.2 it has an even number. Since it is not major, it follows easily that there are two edges of C so that the neighbours of p_1 in C are the ends of these two edges, and therefore P is clear, a contradiction. This proves (2).

Since P is a shortcut, and P has length $k+1$, it follows that $k + 1 < t/4$, and so $t \geq 4k + 5 \geq 13$.

(3) *There are disjoint subpaths Q_1, Q_2 of C, both of length ≤ 2, such that $V(Q_1)$ is the set of neighbours of p_1 in C, and $V(Q_2)$ is the set of neighbours of p_k in C.*

Since $k \geq 2$, it follows that p_1 is not the unique internal vertex of any bad shortcut; and so every two of its neighbours in C are joined by a path in C of length ≤ 2. Since C has length $\geq 13 \geq 7$, there is a path Q_1 in C of length ≤ 2 containing all neighbours of p_1 in C. Choose Q_1 minimal. We claim that p_1 is adjacent to every vertex of Q_1. For the minimality of Q_1 implies that p_1 is adjacent to the ends of Q_1, so we only need check that if Q_1 has length 2 then p_1 is adjacent to its middle vertex. Let Q_1 be c_1-c_2-c_3 say. Then the path c_1-p_1-c_3 is not a bad shortcut, and therefore it has the same γ-parity as the path c_1-c_2-c_3. If the union of these paths is a hole, then this hole has length 4 and γ-parity 0, contradicting that $t \geq 6$. So it is not a hole, and therefore p_1 is adjacent to c_2. This proves that p_1 is adjacent to every vertex of Q_1. Define Q_2 similarly for p_k. Suppose Q_1 meets Q_2. Then $Q_1 \cup Q_2$ is a path of length ≤ 4 with ends $q_1 \in V(Q_1)$ and $q_2 \in V(Q_2)$ say. If $d_C(q_1, q_2) > d_C(u, v)$, then q_1-p_1-$\cdots$-p_k-q_2 is a bad shortcut, worse than P, a contradiction. So $d_C(q_1, q_2) \leq d_C(u, v)$, and since C has length at least 13, and u, v belong to $Q_1 \cup Q_2$, it follows that $u = q_1$ and $v = q_2$; but then P is shallow, since it has length $k + 1 \geq 3$ and $d_C(u, v) \leq 4$, a contradiction. So Q_1, Q_2 are disjoint. This proves (3).

(4) *There are edges between $\{p_2, \dots, p_{k-1}\}$ and $V(C)$.*

For suppose there are no such edges; then the only edges between the interior of P and C are those incident with p_1 or with p_k. Let Q_1, Q_2 be as in (3). Since P is not clear, one of Q_1, Q_2 has even length, and the other is odd; say Q_1 has even length, and Q_2 has length 1. In the subgraph induced on $V(C) \cup \{p_1, \dots, p_k\}$ there is a set of either four or six holes, depending whether p_1 has one or three neighbours in C, such that every edge is in an even number of them (C itself, the two holes that include p_1-$\cdots$-p_k, and either one or three holes of length 3); one of these holes is C, and so by 2.1, one of the others is γ-even and has length $\geq t$. Hence we may assume that there are ends q_i of Q_i ($i = 1, 2$), such that no internal vertex of $C(q_2, q_1)$ belongs to Q_1 or to Q_2, and the hole q_1-p_1-$\cdots$-p_k-q_2-$C(q_2, q_1)$-q_1 is γ-even and has length $\geq t$. It follows that $C(q_2, q_1)$ has length at least $t - k - 1$, and so $C(q_1, q_2)$ has length at most $k + 1$. Consequently $d_C(u, v) \leq k + 1$, with strict inequality unless $u = q_1$ and $v = q_2$; and since P is a shortcut it follows that equality holds and $u = q_1, v = q_2$. But the hole induced on the union of P and $C(v, u)$ is therefore another shortest γ-even hole, contradicting that P is a bad shortcut. This proves (4).

(5) *$C(u,v)$ has length at most $k+3$, and $C(v,u)$ has length at least $3k+2$, and there are no edges between $\{p_1,\ldots,p_k\}$ and the interior of $C(v,u)$.*

For by (4), p_i is adjacent to c_j for some i,j with $2 \le i \le k-1$ and $1 \le j \le t$. Let P' be a path between u, c_j with interior in $\{p_1,\ldots,p_i\}$. Then the interior of P' is a proper subset of that of P, and so by hypothesis, P' is not a bad shortcut. Consequently $d_C(u,c_j) \le i+1$. Similarly $d_C(v,c_j) \le k-i+2$ and so $d_C(u,v) \le (i+1)+(k-i+2) = k+3$. Since $C(u,v)$ is the shorter of the two paths of C between u,v, it follows that $C(u,v)$ has length at most $k+3$. Since C has length $\ge 4k+5$, we deduce that $C(v,u)$ has length at least $3k+2$. That proves the first two assertions of (5). Now suppose that c_j belongs to $C(v,u)$. Since $d_C(u,c_j) \le i+1$ and $C(u,v)$ has length $\ge k+1 > i+1$, it follows that the shortest path of C between u, c_j does not include $C(u,v)$, and therefore is $C(c_j,u)$. We deduce that $C(c_j,u)$ has length at most $i+1$. Similarly $C(v,c_j)$ has length at most $k-i+2$, and so $C(v,u)$ has length at most the sum of these, that is, at most $k+3$, a contradiction. So there are no edges between $\{p_2,\ldots,p_{k-1}\}$ and $C(v,u)$. Finally, suppose that there is an edge between say p_1 and the interior of $C(v,u)$. Consequently there is a vertex $q \in V(Q_1)$, not in $C(u,v)$ and adjacent to u. But then the path $q\text{-}p_1\text{-}\cdots\text{-}p_k\text{-}v$ is a bad shortcut, worse than P, a contradiction. This proves (5).

Since P is not shallow, it follows from (5) that $C(u,v)$ has length $> k+2$. Since $C(u,v)$ has length at most $k+3$, it has length exactly $k+3$. Let $u = c_1, v = c_{k+4}$ say. By (4), p_i is adjacent to c_j for some i,j with $2 \le i \le k-1$ and $1 \le j \le k+4$. As we saw above, $d_C(u,c_j) \le i+1$, that is, $j \le i+2$, and $d_C(v,c_j) \le k-i+2$, that is, $j \ge i+2$. So equality holds, and in particular, $c_1\text{-}p_1\text{-}\cdots\text{-}p_i\text{-}c_j$ is a path (P' say) with the same length as $C(c_1,c_j)$. Since P' is not a worse shortcut than P, it is not a bad shortcut at all, and so P' has the same γ-parity as $C(c_1,c_j)$. Similarly, $c_j\text{-}p_i\text{-}\cdots\text{-}p_k\text{-}c_{k+4}$ is a path (P'' say), and P'' has the same γ-parity as $C(c_j,c_{k+4})$. It follows that P has the same γ-parity as $C(u,v)$; but then the union of P and $C(v,u)$ is a γ-even hole of length $< t$, a contradiction. This proves 3.1. ∎

4 Eliminating clear shortcuts

In this section we give another cleaning subroutine, an extension of 2.5; this time, we will generate polynomially many subsets such that for one of them (say X), no vertex of C is in X, and if in $G \setminus X$ there is still a bad shortcut across C, then the worst is shallow. We need the following two lemmas.

4.1 *Let (G,γ) be a signed graph, and let C be a shortest γ-even hole in G. Let P be a worst shortcut, and assume that P is clear. Let P be $u\text{-}p_1\text{-}\cdots\text{-}p_k\text{-}v$, and let R be a minimum length path in G between p_1, p_k. Then the interior of R is separate from $V(C)$, and $u\text{-}p_1\text{-}R\text{-}p_k\text{-}v$ is another worst shortcut, again clear.*

Proof. If $k \le 2$ the claim is trivial, so we assume $k \ge 3$. Let C have vertices $c_1\text{-}\cdots\text{-}c_t$ in order, where the neighbours of p_1 in C are $c_1,\ldots,c_h$ and those of p_k are $c_i,\ldots,c_j$. So $1 \le h < i \le j \le t$. Since P is clear, $d_C(u,v) \ge k+3$. Let P' be a path between u,v with interior in $V(R)$. Since R has length $\le k-1$ (because $p_1\text{-}\cdots\text{-}p_k$ is a path between p_1,p_k), it follows that P' has length at most $k+1$, and therefore P' is a bad shortcut. Since its length is at most that of P, and P is a shortest bad shortcut, it follows that P' has length $k+1$ and is another shortest bad shortcut, and has interior $V(R)$. Since it has the same ends as P, P' is another worst shortcut. By 3.1, either P' has only one internal vertex and that vertex is major, or P' is clear, or P' is shallow. The first is impossible since there are no major vertices for C (because P is a shortest bad shortcut and has length > 2). The second implies the desired result; and the third is impossible since P is deep. This proves 4.1. ∎

Let P be a clear shortcut across C, with attachment paths Q_1, Q_2. If P is of even type, we define the *anchors* of P to be the middle vertices of Q_1, Q_2, and if it is of odd type, we define its anchors to be the edges of Q_1, Q_2.

4.2 *Let (G, γ) be a signed graph, and let C be a shortest γ-even hole in G. Let P, Q be clear shortcuts, both of the same type. Then either*

- *the interiors of P, Q are not separate, or*
- *P, Q share an anchor.*

Proof. Let P,Q have interiors $p_1\text{-}\cdots\text{-}p_k$ and $q_1\text{-}\cdots\text{-}q_l$ respectively, and assume that neither of the outcomes hold. Suppose first that there is an edge of C which is in one of the attachment paths for P and in one of them for Q. Since P, Q do not share an anchor, it follows that they have even type. The two attachment paths which share an edge therefore both have length 2, and so $k, l \geq 2$; and since P, Q do not share an anchor, we may assume that p_1 is adjacent to c_1, c_2, c_3 and q_1 is adjacent to c_2, c_3, c_4, where C has vertices $c_1, \ldots, c_t$ as usual. Choose i with $1 \leq i \leq t$ minimum such that p_k is adjacent to c_i. The hole $c_2\text{-}p_1\text{-}\cdots\text{-}p_k\text{-}c_i\text{-}C(c_4, c_i)\text{-}q_1\text{-}c_2$ is shorter than C (since P is clear and therefore deep), and so this hole is γ-odd; but c_3 has four neighbours in it, contrary to 2.2.

It follows that the attachment paths for P, Q are edge-disjoint. Since P, Q share no anchor, it follows that any vertex that is in both an attachment path for P and one for Q is an end of both paths. Let C have vertices $c_1\text{-}\cdots\text{-}c_t$ in order. We may assume that the attachment paths of P are $C(c_1, c_g)$ and $C(c_{i_1}, c_{i_2})$, where p_1 is adjacent to c_1; and those for Q are $C(c_{h_1}, c_{h_2})$ and $C(c_{j_1}, c_{j_2})$, where q_1 is adjacent to c_{h_1}. Suppose first that they do not alternate; say

$$1 \leq g \leq h_1 \leq h_2 \leq j_1 \leq j_2 \leq i_1 \leq i_2 \leq t.$$

Let C_1, C_2, C_3 be the following three holes:

$$c_g\text{-}C(c_g, c_{i_1})\text{-}c_{i_1}\text{-}p_k\text{-}\cdots\text{-}p_1\text{-}c_g$$
$$c_{h_1}\text{-}q_1\text{-}\cdots\text{-}q_l\text{-}c_{j_2}\text{-}C(c_{j_2}, c_{h_1})\text{-}c_{h_1}$$
$$c_{h_1}\text{-}q_1\text{-}\cdots\text{-}q_l\text{-}c_{j_2}\text{-}C(c_{j_2}, c_{i_1})\text{-}c_{i_1}\text{-}p_k\text{-}\cdots\text{-}p_1\text{-}c_g\text{-}C(c_g, c_{h_1})\text{-}c_{h_1}.$$

Then they are all of length $< t$ and yet every edge is in an even number of C, C_1, C_2, C_3, contrary to 2.1. Next assume the paths do alternate; that is, say,

$$1 \leq g \leq h_1 \leq h_2 \leq i_1 \leq i_2 \leq j_1 \leq j_2 \leq t.$$

Let C_1, C_2 be the following two holes:

$$c_g\text{-}C(c_g, c_{h_1})\text{-}c_{h_1}\text{-}q_1\text{-}\cdots\text{-}q_l\text{-}c_{j_1}\text{-}C(c_{i_2}, c_{j_1})\text{-}c_{i_2}\text{-}p_k\text{-}\cdots\text{-}p_1\text{-}c_g$$
$$c_{h_2}\text{-}C(c_{h_2}, c_{i_1})\text{-}c_{i_1}\text{-}p_k\text{-}\cdots\text{-}p_1\text{-}c_1\text{-}C(c_{j_2}, c_1)\text{-}c_{j_2}\text{-}q_l\text{-}\cdots\text{-}q_1\text{-}c_{h_2}.$$

The sum of their lengths is at most $2k + 2l + 4 + t$, and since $k + 1, l + 1 < t/4$, the sum of their lengths is less than $2t$. Consequently one of them has length less than t and from the symmetry we may assume it is C_1. Let C_3, C_4 be the following two holes:

$$c_1\text{-}p_1\text{-}\cdots\text{-}p_k\text{-}c_{i_2}\text{-}C(c_{i_2}, c_1)\text{-}c_1$$
$$c_{h_1}\text{-}q_1\text{-}\cdots\text{-}q_l\text{-}c_{j_2}\text{-}C(c_{j_2}, c_{h_1})\text{-}c_{h_1},$$

and append to the list C_1, C_3, C_4 all the triangles within $\{p_1, c_1, \ldots, c_g\}$ and all the triangles within $\{q_l, c_{j_1}, \ldots, c_{j_2}\}$. The total number of triangles we append is even since P, Q have the same type. Every edge is in an even number of the members of the expanded list, and yet they all have length $< t$, contrary to 2.1. This proves 4.2. ∎

4.3 *There is an algorithm with the following specifications:*

- **Input** *A signed graph* (G, γ).
- **Output** *A sequence of subsets* $X_1, \ldots, X_r$ *of* $V(G)$ *with* $r \le |V(G)|^8$, *such that, if* C *is a smallest* γ*-even hole in* G, *and there is a worst shortcut across* C, *and it is clear and of even type, then there is a member* X *of the output list such that* $X \cap V(C) = \emptyset$ *and there is no clear shortcut of even type across* C *in* $G \setminus X$.
- **Running time** $O(|V(G)|^8)$.

Proof. Here is the algorithm. For every quadruple v_1, v_2, v_3, v_4 of vertices, such that v_1, v_2 are in the same component of G and v_1v_3, v_2v_4 are edges, we proceed as follows. Let Z be the set of all vertices adjacent to at least one of v_3, v_4; find a shortest path R between v_1, v_2 in G; and find the set Y of all vertices different from v_3, v_4, that are either in R or with a neighbour in R; remove at most four vertices from $Y \cup Z$, in all possible ways; and output the sets we generate. Finally, output $\emptyset$.

The algorithm outputs at most $|V(G)|^8$ subsets, and its running time is as claimed. We must prove that the output sequence has the required property. So let C be a shortest γ-even hole in G, and let P be a worst shortcut across C, say $u\text{-}p_1\text{-}\cdots\text{-}p_k\text{-}v$, and assume that it is clear and of even type. Let q_1, q_2 be the anchors of P, where p_1 is adjacent to q_1. Let R be the shortest path between p_1, p_k chosen by the algorithm when it examines the quadruple p_1, p_k, q_1, q_2. By 4.1, $u\text{-}p_1\text{-}R\text{-}p_k\text{-}v$ is another worst shortcut, and it is again clear. Let Y be the set of all vertices different from q_1, q_2 that are either in R or with a neighbour in R. Let Z be the set of all vertices in G adjacent to one of q_1, q_2. So there are at most four vertices of $Y \cup Z$ that belong to C; let X be the members of $Y \cup Z$ that are not in C. So X is a member of the list output by the algorithm. Clearly it is disjoint from C, and by 4.2 it meets all clear shortcuts of even type. This proves 4.3. ∎

We need a similar subroutine to clean shortcuts of odd type, as follows.

4.4 *There is an algorithm with the following specifications:*

- **Input** *A signed graph* (G, γ).
- **Output** *A sequence of subsets* $X_1, \ldots, X_r$ *of* $V(G)$ *with* $r \le |V(G)|^6$, *such that, if* C *is a smallest* γ*-even hole in* G, *and there is a worst shortcut across* C, *and it is is clear and of odd type, then there is a member* X *of the output list such that* $X \cap V(C) = \emptyset$ *and there is no clear shortcut of odd type across* C *in* $G \setminus X$.

- **Running time** $O(|V(G)|^6)$.

Proof. Here is the algorithm. For every edge uv, find the set of all vertices adjacent to both u, v, say $Z(uv)$. For every pair of vertices (u, v) in the same component, find a shortest path $R(u, v)$ between u, v, and find the set $Y(u, v)$ of all vertices that are either in $R(u, v)$ or have a neighbour in it. For each 6-tuple $v_1, \ldots, v_6$ with v_1, v_2 in the same component, such that v_3v_4 and v_5v_6 are edges, output the set

$$(Y(v_1, v_2) \cup Z(v_3v_4) \cup Z(v_5v_6)) \setminus \{v_3, v_4, v_5, v_6\}.$$

Finally, output $\emptyset$.

Again we must show that the output has the desired property. So let C be a shortest γ-even hole in G, and let P be a worst shortcut across C, say u-p_1-$\cdots$-p_k-v, and assume that P is clear and of odd type. Let v_3v_4, v_5v_6 be the anchors of P, where p_1 is adjacent to v_3. Let

$$X = (Y(p_1, p_k) \cup Z(v_3v_4) \cup Z(v_5v_6)) \setminus \{v_3, v_4, v_5, v_6\}.$$

So Y is one of the sets output by the algorithm; we claim it has the property we need. No vertex of C belongs to $Z(v_3v_4) \cup Z(v_5v_6)$, since v_3v_4, v_5v_6 are disjoint edges of C. Moreover, by 4.1, u-p_1-$R(p_1, p_k)$-p_k-v is another worst shortcut, and it is again clear and of odd type. Consequently, no vertex of $Y(p_1, p_k)$ belongs to C except $v_3, \ldots, v_6$. Moreover, every clear shortcut of odd type contains a vertex of X, by 4.2. This proves 4.4. ∎

Next we combine our three cleaning subroutines into one, as follows.

4.5 *There is an algorithm with the following specifications:*

- **Input** *A signed graph* (G, γ).
- **Output** *A sequence of subsets* $X_1, \ldots, X_r$ *of* $V(G)$ *with* $r \leq 2|V(G)|^{23}$, *such that, if* C *is a smallest* γ*-even hole in* G, *then there is a member* X *of the output list such that* $X \cap V(C) = \emptyset$ *and in* $G \setminus X$, *every worst shortcut across* C *is shallow.*
- **Running time** $O(|V(G)|^{23})$.

Proof. The algorithm is in three phases, as follows. First we run 2.5 on G; let $X_1, \ldots, X_a$ be the subsets we generate. In phase 2 we examine all the graphs $G \setminus X_i$ in turn. Fix i with $1 \leq i \leq a$. We run 4.3 on $G \setminus X_i$ (let $Y_1, \ldots, Y_{b_i}$ be the sets we generate), and then run 4.4 on each of the graphs $G \setminus (X_i \cup Y_j)$ for $1 \leq j \leq b_i$; and for each j, and for each member of the output list from 4.4 applied to $G \setminus (X_i \cup Y_j)$, we output its union with $X_i \cup Y_j$. We repeat this for all i.

In phase 3 we repeat phase 2, this time running 4.4 before 4.3. So again, for $1 \leq i \leq a$, we run 4.4 on $G \setminus X_i$, generating $Y_1, \ldots, Y_{b_i}$ say; run 4.3 on each of the graphs $G \setminus (X_i \cup Y_j)$; and output the union of each of the output sets with $X_i \cup Y_j$. That completes the description of the algorithm.

The running time is evidently as claimed, and the number of output sets is as claimed. We must check that the output has the desired property. So let C be a smallest γ-even hole in G. There is one of the X_i's produced by 2.5 such that X_i is disjoint from C and contains all major vertices. Suppose that in $G \setminus X_i$ there is a worst shortcut across C, and it is clear and of even type. When we run 4.3 in phase 2 on $G \setminus X_i$, one of the sets Y_j we output has the property that Y_j is disjoint from C and meets all clear shortcuts of even type across C in $G \setminus (X_i \cup Y_j)$. So in $G \setminus (X_i \cup Y_j)$, every worst shortcut is either clear and of odd type, or shallow, and there is no clear shortcut of even type. Now $X_i \cup Y_j$ belongs to the output list of 4.5 (because $\emptyset$ is always one of the sets in the output of 4.4), so if every worst shortcut in $G \setminus (X_i \cup Y_j)$ is shallow, then we have the desired property. We may therefore assume that there is a worst shortcut in $G \setminus (X_i \cup Y_j)$ that is clear and of odd type. When we apply 4.4 to $G \setminus (X_i \cup Y_j)$ in phase 2, we output a set Z_k that is disjoint from C and meets every clear shortcut of odd type. It follows that in $G \setminus (X_i \cup Y_j \cup Z_k)$, there is no major vertex, and no clear shortcut. In particular, by 3.1, every worst shortcut is shallow. Consequently in this case we have been successful.

We may therefore assume that in $G \setminus X_i$ there is no worst shortcut across C that is clear and of even type. If there is one that is clear and of odd type, then the argument is similar, using phase 3 instead of phase 2. Finally, if there is no worst shortcut that is clear at all, then we have succeeded since X_i is one of the outputs of 4.5. This proves 4.5. ∎

5 Detecting shallow shortcuts

That concludes the "cleaning" part of the algorithm. So far, we have not detected any γ-even holes at all; in fact the algorithm has not even read the γ-values on the edges of G. In this section we give an algorithm that attempts to find a γ-even hole. It might not succeed, even if such a hole is present, but it is guaranteed to succeed if there is a shortest γ-even hole C such that some worst shortcut across C is shallow. We need the following lemma.

5.1 *Let C be a shortest γ-even hole in G, with vertices $c_1, \dots, c_t$ in order. Let P be a worst shortcut across C, and assume P is shallow. Let P be c_1-p_1-$\cdots$-p_k-c_i say, and assume that $C(c_1, c_i)$ has length at most that of $C(c_i, c_1)$. (So $i = k+2$ or $k+3$.) Then*

- *every path between c_1, c_i in G has length $\geq k+1$*
- *for every path Q of length $k+1$ between c_1, c_i, the interiors of Q and of $C(c_i, c_1)$ are separate*
- *every path between c_2, c_{i-1} has length at least $i-3$*
- *for every path Q of length $i-3$ between c_2, c_{i-1}, the interiors of Q and of $C(c_{i-1}, c_2)$ are separate.*

Proof. Since P is a shortcut, $d_C(c_1, c_i) \geq k+1$. Since $C(c_1, c_i)$ has length at most that of $C(c_i, c_1)$, it follows that $d_C(c_1, c_i) = i-1$; and since P is shallow, we deduce that $i = k+2$ or $k+3$.

Suppose there is a major vertex for C, say w. Then $k = 1$, and $t \geq 4(k+1)+1 = 9$. Since w has at least four neighbours in C, and three are pairwise nonadjacent, it follows that w has two neighbours u', v' with $d_C(u', v') \geq 4$, contrary to the choice of the ends of P. So there is no major vertex for C.

If Q is a path between c_1, c_i of length $< k+1$, then it is a bad shortcut, which is impossible since P is a worst shortcut. This proves the first assertion.

For the second, assume that Q is a path of length $k+1$ between c_1, c_i. Then Q is a shortcut, and if it is good then the claim follows, so we assume it is bad. Since it has the same length and ends as P, it is another worst shortcut. In particular, since $d_C(c_t, c_i) > d_C(c_1, c_i)$, it follows that no internal vertex of Q is adjacent to p_1, and similarly no internal vertex of Q is adjacent to c_{i+1}. Thus if Q is clear then the second assertion follows. But it also follows trivially if Q is shallow, and since there are no major vertices, the result follows from 3.1 applied to Q. This proves the second assertion.

For the third assertion, suppose that Q is a path between c_2, c_{i-1} of length $< i-3$. Then Q is a bad shortcut strictly shorter than P, a contradiction. This proves the third assertion.

For the fourth, suppose that Q is a path between c_2, c_{i-1} of length $i-3$. So Q is a shortcut; and since P is a worst shortcut, it follows that Q is good, and in particular the fourth assertion follows. This proves 5.1. ∎

For the main algorithm of this section, we need the following:

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

- **Input** *A signed graph* (G, γ).
- **Output** *For each pair of vertices* u, v, *it computes* $d_G(u, v)$; *and for* $z = 0, 1$, *it finds a path of* G *between* u, v *with length* $d_G(u, v)$ *and* γ*-parity* z *if one exists.*
- **Running time** $O(|V(G)|^3)$.

Proof. We can compute all the numbers $d_G(u, v)$ in time $O(|V(G)|^3)$, using Dijkstra's algorithm repeatedly. Now make a new graph H as follows. To each vertex x of G there correspond two vertices x_1, x_2 of H; and to each edge $e = xy$ of G there correspond two edges of H, x_1y_1, x_2y_2 if $\gamma(e) = 0$, and x_1y_2, x_2y_1 if $\gamma(e) = 1$. Now for all u, v in $V(G)$, we test whether there is a path in H of length $d_G(u, v)$ between u_1, v_1 and also test for such a path between u_1, v_2. These paths exist if and only if there is a path in G between u, v of length $d_G(u, v)$ and γ-parity 0 and 1 respectively. For the pairs of H where such a path exists, we find such a path and output the corresponding path of G. This proves 5.2. ∎

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

- **Input** *A signed graph* (G, γ).
- **Output** *Either:*
  - *a* γ*-even hole in* G, *or*
  - *a determination that there is no shortest* γ*-even hole in* G *such that some worst shortcut is shallow.*
- **Running time** $O(|V(G)|^6)$.

Proof. The algorithm is in three phases.

First we run 5.2; for each pair u, v of vertices, we find $d_G(u, v)$, and find a path of this length between u, v and of γ-parity z, for $z = 0, 1$ if such a path exists. Call the path we find $P_z(u, v)$.

In phase 2, for every pair of distinct vertices u, v such that $P_0(u, v)$ and $P_1(u, v)$ both exist, let W be the union of the interiors of these two paths; we test whether there is a path Q between u, v with interior separate from W. If there is such a path, then Q has the same γ-parity as one of $P_0(u, v), P_1(u, v)$, and so their union is a γ-even hole, and we output it and stop.

In phase 3, we examine all quadruples u, v, w, x of distinct vertices and $y, z \in \{0, 1\}$ such that

- u, v are adjacent, and so are w, x; let $uv = e$ and $wx = f$ say
- $y + z + \gamma(e) + \gamma(f) = 1$
- $d_G(u, x) = d_G(v, w) - 1$
- $P_y(x, u)$ and $P_z(w, v)$ both exist, and
- v and w have no neighbour in $P_y(x, u)$ except u and x respectively (and in particular, v, w are nonadjacent).

For every such choice of u, v, w, x, y, z, let

$$W = V(P_y(x, u)) \cup V(P_z(w, v)) \setminus \{w, v\}.$$

We test whether there is a path Q in G between v, w with interior separate from W. If Q exists, then it has the same γ-parity as one of the paths w-x-$P_y(x, u)$-u-v, $P_z(w, v)$; and so their union is a γ-even hole and and we output it and stop. If we do not find such a hole, we report that there is no shortest γ-even hole in G such that some worst shortcut is shallow.

This evidently has running time $O(|V(G)|^6)$, but we have to check that the output is correct. Suppose then that there is a shortest γ-even hole C, and some worst shortcut P across C is shallow. Let C have vertices $c_1, \dots, c_t$ in order, and let P have vertices c_1-p_1-$\cdots$-p_k-c_i, where $i = k + 2$ or $k + 3$ and there are no edges between $\{p_1, \dots, p_k\}$ and the interior of $C(c_i, c_1)$. Since P is a bad shortcut, the hole formed by the union of P and $C(c_i, c_1)$ is γ-odd, and so P, $C(c_1, c_i)$ have opposite γ-parity. By the first assertion of 5.1, $d_G(c_1, c_i) = k + 1$.

Suppose first that $i = k + 2$. Since $P, C(c_1, c_i)$ have opposite γ-parity and both have length $d_G(c_1, c_i)$, it follows that the algorithm will find two paths $P_0(c_1, c_i), P_1(c_1, c_i)$. Let W denote the union of the interiors of these two paths. Now there is a path Q between c_1, c_i with interior separate from W (for $C(c_i, c_1)$ is such a path, by the second assertion of 5.1); and therefore the algorithm will detect such a path and output a γ-even hole. So in this case the output is correct.

Now assume that $i = k + 3$. Then $d_G(c_{i-1}, c_2) = i - 3 = k$, by the third assertion of 5.1. Let y, z be the γ-parities of $C(c_2, c_{i-1}), P$ respectively; then $y + z + \gamma(c_1c_2) + \gamma(c_{i-1}c_i) = 1$, since $P, C(c_1, c_i)$ have opposite γ-parity. The algorithm will choose paths $P_y(c_{i-1}, c_2)$ and $P_z(c_i, c_1)$, since they both exist. Let

$$W = V(P_y(c_{i-1}, c_2)) \cup V(P_z(c_i, c_1)) \setminus \{c_i, c_1\}.$$

There is a path Q between c_1, c_i with interior separate from W, because $C(c_i, c_1)$ is such a path, by the second and fourth assertions of 5.1. Consequently the algorithm will detect such a path and output a γ-even hole. So in all cases the output is correct. This proves 5.3. ∎

6 The complete algorithm

Before we can put these pieces together, we need one more subroutine; this one will detect a γ-even hole in the case when there is a shortest γ-even hole with no bad shortcut at all.

6.1 *There is an algorithm with the following specifications:*

- **Input** *A signed graph* (G, γ)*, such that there is no* γ*-even hole in* G *with length at most* 12.
- **Output** *Either:*
  - *a* γ*-even hole in* G*, or*
  - *a determination that there is no shortest* γ*-even hole in* G *such that every shortcut across it is good.*
- **Running time** $O(|V(G)|^8)$.

Proof. The algorithm is as follows. For every pair of vertices u, v in the same component of G, find $d_G(u, v)$, and find a path $P(u, v)$ between u, v of this minimum length. Next, for every 8-tuple $v_1, \ldots, v_8$ of distinct vertices in the same component of G, test whether the union of the eight paths $P(v_1, v_2), P(v_2, v_3), \ldots, P(v_8, v_1)$ is a γ-even hole. If so, output this hole and stop. If we do not find such a hole, output that there is no shortest γ-even hole in G such that every shortcut across it is good.

This has running time $O(|V(G)|^{10})$, but with a little more care we can get it down to $O(|V(G)|^8)$, as follows. First, for each pair u, v we find the γ-parity of $P(u, v)$. Then, we compute the triples u, v, w such that $P(u, v) \cup P(v, w)$ is a path from u to w; and we compute the quadruples u, v, w, x such that $P(u, v)$ and $P(w, x)$ are disjoint and there is no edge between $P(u, v)$ and $P(w, x)$ except possibly between their ends. And given all that information, we examine each 8-tuple $v_1, \ldots, v_8$ as before; but now we can process each one in constant time.

We must also check that the output is correct. Suppose then that there is a shortest γ-even hole C, and every shortcut across C is good. Let C have length t say. By hypothesis, $t \geq 13$. Let $t = 8a + b$, where a, b are integers and $0 \leq b \leq 7$. Choose eight distinct vertices $v_1, \ldots, v_8$ in clockwise order on C; let $C_1 = C(v_1, v_2), C_2 = C(v_2, v_3)$ and so on. Choose $v_1, \ldots, v_8$ so they are as evenly spaced as possible; more precisely,

- each C_i has length a or $a + 1$
- if $b \leq 4$ then no two consecutive of $C_1, \ldots, C_8$ both have length $a + 1$.

(1) *The union of any two consecutive of* $C_1, \ldots, C_8$ *has length* $< t/4 + 1$.

For if $b = 0$ then this is clear; if $1 \leq b \leq 4$ then $C_1 \cup C_2$ (say) has length at most $2a+1 < (8a+b)/4+1$ since $b > 0$; and if $b > 4$ then $C_1 \cup C_2$ has length at most $2a + 2 < (8a + b)/4 + 1$ since $b > 4$. This proves (1).

(2) $P(v_1, v_2)$ *has the same length and* γ*-parity as* $C(v_1, v_2)$*, and* $v_3, \ldots, v_8$ *do not belong to it and have no neighbours in its interior.*

For let $C(v_1, v_2)$ have length k say. Then $P(v_1, v_2)$ exists and has length $\leq k$; and since there is no bad shortcut over C, it follows that $P(v_1, v_2)$ has length k and its interior is separate from the interior of $C(v_2, v_1)$; and $P(v_1, v_2)$ has the same γ-parity as $C(v_1, v_2)$. This proves (2).

(3) $P(v_1, v_2) \cup P(v_2, v_3)$ *is a path from* v_1 *to* v_3.

For suppose not; then there is a path P between v_1, v_3 whose length is strictly less than the sum of the lengths of $C(v_1, v_2), C(v_2, v_3)$. But by (1) the latter is $< t/4 + 1 \leq t/2$, and therefore equals $d_C(v_1, v_3)$, and so P has length at most $d_C(v_1, v_3) - 1 < t/4$. In particular, P is a bad shortcut across C, a contradiction. This proves (3).

(4) *$P(v_1, v_2), P(v_3, v_4)$ are disjoint, and there is no edge between them, except possibly the edge v_2v_3 if C_2 has length 1.*

For suppose not. If $P(v_1, v_2), P(v_3, v_4)$ both have length at most 2, then there is a path of length ≤ 3 between v_1, v_4, which is therefore a bad shortcut (since C has length ≥ 13), a contradiction. So we may assume that one of $P(v_1, v_2), P(v_3, v_4)$ has length ≥ 3. Consequently one of C_1, C_3 has length ≥ 3, and so $a \geq 2$ and $t \geq 17$. There are two paths P, Q, such that P is between v_1, v_3, and Q is between v_2, v_4, sharing at most two vertices, and both with interior in W, where W is the union of the interiors of $P(v_1, v_2)$ and $P(v_3, v_4)$. Consequently, the sum of the lengths of P and Q is at most $d_C(v_1, v_2) + d_C(v_3, v_4) + 2$. Since P is not a bad shortcut, it has length at least $d_C(v_1, v_3)$ as before, and similarly Q has length at least $d_C(v_2, v_4)$; and so

$$d_C(v_1, v_2) + d_C(v_3, v_4) + 2 \geq d_C(v_1, v_3) + d_C(v_2, v_4),$$

which implies that v_2v_3 is an edge, a contradiction since $a \geq 2$. This proves (4).

(5) *For $i = 4, 5, 6$, $V(P(v_1, v_2)), V(P(v_i, v_{i+1}))$ are separate.*

For suppose not. As in the previous case, it follows that $a \geq 2$ and $t \geq 17$. There are two paths P, Q, such that P is between v_1, v_i, and Q is between v_2, v_{i+1}, sharing at most two vertices, and both with interior in W, where W is the union of the interiors of $P(v_1, v_2)$ and $P(v_i, v_{i+1})$. The sum of the lengths of P and Q is at most $d_C(v_1, v_2) + d_C(v_i, v_{i+1}) + 2$. Since P is not a bad shortcut, it has length at least $\min(t/4, d_C(v_1, v_i))$, and in particular, P has length at least $a + 3$. So does Q, and hence $2(a+3) \leq d_C(v_1, v_2) + d_C(v_i, v_{i+1}) + 2$, which is impossible since $d_C(v_1, v_2), d_C(v_i, v_{i+1}) \leq a+1$. This proves (5).

From (2)-(5), the union of the eight paths $P(v_1, v_2), \ldots, P(v_8, v_1)$ is a hole, and from (2), this hole has the same γ-parity as C, and is therefore γ-even. Consequently, in this case the algorithm correctly outputs a γ-even hole. This proves 6.1. ∎

Now let us assemble the pieces, in the following, the main result of the paper.

6.2 *There is an algorithm with the following specifications:*

- **Input** *A signed graph (G, γ).*
- **Output** *Either:*
  - *a γ-even hole in G, or*
  - *a determination that there is no γ-even hole in G.*
- **Running time** $O(|V(G)|^{31})$.

Proof. Here is the algorithm. First we test whether there is a γ-even hole with length at most 12, by examining all 12-tuples of vertices. If we find one we stop. Otherwise, we run 4.5, and generate a sequence $X_1, \ldots, X_r$ of subsets of $V(G)$. For $i = 1, \ldots, r$, we run both 5.3 and 6.1 on $G \setminus X_i$. If at any stage we find a γ-even hole, we output it and stop; and otherwise we output that there is no such hole.

This has running time $O(|V(G)|^{31})$. To see that the output is correct, suppose that G has a γ-even hole, and let C be a shortest such hole. Then there is some X_i such that X_i is disjoint from C, and every worst shortcut across C in $G \setminus X_i$ is shallow. If there is a worst shortcut across C, then when we run 5.3 on $G \setminus X_i$ we find a γ-even hole; and if there is none, then every shortcut across C is good, and we detect a γ-even hole when we run 6.1 on $G \setminus X_i$. In either case we detect a γ-even hole, and so the output is correct. This proves 6.2. ∎

7 Variations and refinements

The running time $O(|V(G)|^{31})$ is a little embarrassing. Our excuse is that we decided to present the simplest polynomial-time algorithm that we could find; there are places where we could have made it run faster with a little more work. But not that much faster, in fact; we can get it down to about $O(|V(G)|^{15})$, but not much less. In this section we sketch how to do so.

The algorithm presented here was constructed by adapting the algorithm in [1, 2] to test for Bergeness. The main difference between the algorithms for the even hole problem and for the odd hole problem is that when we are looking at the shortest odd hole, we can prove that if there is no major vertex then there is no bad shortcut. For the shortest even hole, we can't prove this; we can only arrange it by cleaning. In the first case, we know something about *every* shortest odd hole, and in the second case we only know it about *some* shortest even hole (assuming there is one); and that makes a significant difference. For instance, suppose we have arranged that some shortest even hole C has no bad shortcuts, by cleaning. When we reroute C along a good shortcut, we get another shortest even hole, and it would be helpful if this one was as good as C; but it might not be, because the second hole might have bad shortcuts. This is the reason why we have to guess eight vertices in 6.1; in the analogous result of [1] it was enough to guess three vertices of the shortest odd hole.

The reason that we could prove there were no bad shortcuts in [1] was that we were sure that the graph contained no pyramids. (A *pyramid* is an induced subgraph consisting of three pairwise adjacent vertices b_1, b_2, b_3, a fourth vertex a adjacent to at most one of b_1, b_2, b_3, and three paths from a to b_1, b_2, b_3 respectively, disjoint except for a, and with no edges between them except those already specified.) Any graph with a pyramid certainly contains an odd hole, and a key idea of that paper was to test for the presence of pyramids first, before we do any cleaning; if we find one we are done, and if we can show that there are none then that greatly simplifies the remainder of the problem.

The analogous operation in our context would be to first test for the presence of thetas and prisms. The nice thing about searching for pyramids is that if there is one, there is a smallest one, and guessing a few important vertices of the smallest pyramid allows one essentially to reconstruct the remainder, by taking shortest paths. One might hope that something similar would work in the even hole case, looking for a smallest theta or prism.

For thetas it does work (although it is nasty — there is one case that doesn't work as one would wish, and we have to start looking for the smallest subgraph that is either a theta or a kind of extended pyramid, not worth describing fully here). Having guessed a few critical vertices, we can reconstruct the remainder, just like for pyramids, in time $O(|V(G)|^{12})$. Henceforth we can therefore assume that G contains no theta. We can (and need to) handle a few related subgraphs similarly. For instance, if there is a hole and a vertex with exactly four neighbours in the hole, three of which are consecutive, then we can detect it by a similar method, and deduce that G has an even hole. Similarly, if there is a hole and a vertex with exactly six neighbours on it, falling into two subpaths each of length 3, we can do the same. So we can assume that G contains neither of these subgraphs. This can all be done in time $O(|V(G)|^{12})$.

What we would really like to do is to test for prisms. The method used above for thetas does not seem to work for prisms directly, however; the only way we can make it work is by combining it with cleaning. First we guess a few important vertices of the smallest prism; then clean the major vertices from it (where "major" means a vertex with three nonadjacent neighbours in the prism) by a version of 2.5 for prisms; and then we can use shortest path methods to reconstruct the prism. Altogether it takes us time $O(|V(G)|^{15})$ to detect prisms.

Henceforth then we can assume that there are no prisms in G either, and that makes things much easier. Now we use 2.5 to clean a shortest even hole; and go directly to 5.3 and 6.1, since there cannot be any clear shortcut, and so we don't have to clean them away. And there is now a faster version of 6.1, in time $O(|V(G)|^{5})$, because now we can prove that if you have a clean shortest even hole, and reroute it along a good shortcut, then you produce another clean shortest even hole. Putting all these pieces together gives an algorithm to detect even holes, with running time $O(|V(G)|^{15})$. But the details are quite complicated and messy, and it does not seem worthwhile to explain them any further here.

References

[1] M. Chudnovsky and P. Seymour, "Recognizing Berge graphs", manuscript, November 2002.

[2] M. Chudnovsky, G. Cornuéjols, X. Liu, P. Seymour and K. Vušković, "Cleaning for Bergeness", manuscript, November 2002.

[3] M. Conforti, G. Cornuéjols, A. Kapoor and K. Vušković, "Even-hole-free graphs Part I: decomposition theorem", *J. Graph Theory* 39 (2002), 6-49.

[4] M. Conforti, G. Cornuéjols, A. Kapoor and K. Vušković, "Even-hole-free graphs Part II: recognition algorithm", *J. Graph Theory* 40 (2002), 238-266.