

On the odd-minor variant of Hadwiger's conjecture

Jim Geelen*, Bert Gerards†, Bruce Reed‡, Paul Seymour§ Adrian Vetta¶

???. revised May 4, 2006

*Department of Combinatorics and Optimization, University of Waterloo. Email: jfgeelen@uwaterloo.ca

†Centrum voor Wiskunde en Informatica, and Eindhoven University of Technology. Email: Bert.Gerards@cwi.nl

‡School of Computer Science, McGill University. Email: breed@cs.mcgill.ca

§Department of Mathematics, Princeton University. Email: pds@math.princeton.edu

¶Department of Mathematics and Statistics, and School of Computer Science, McGill University. Email: vetta@math.mcgill.ca

Abstract

A K_l -*expansion* consists of l vertex-disjoint trees, every two of which are joined by an edge. We call such an expansion *odd* if its vertices can be two-coloured so that the edges of the trees are bichromatic but the edges between trees are monochromatic. We show that, for every l , if a graph contains no odd K_l -expansion then its chromatic number is $O(l\sqrt{\log l})$. In doing so, we obtain a characterization of graphs which contain no odd K_l -expansion which is of independent interest. We also prove that given a graph and a subset S of its vertex set, either there are k vertex-disjoint odd paths with endpoints in S , or there is a set X of at most $2k - 2$ vertices such that every odd path with both ends in S contains a vertex in X . Finally, we discuss the algorithmic implications of these results.

1 The Main Results

All graphs in this paper are finite, and have no loops or parallel edges. Let H, G be graphs. An H -expansion in G is a function η with domain $V(H) \cup E(H)$, satisfying:

- for each $v \in V(H)$, $\eta(v)$ is a subgraph of G that is a tree, and the trees $\eta(v)$ ($v \in V(H)$) are pairwise vertex-disjoint
- for each edge $e = uv$ of H , $\eta(e)$ is an edge $f \in E(G)$, such that f is incident in G with a vertex of $\eta(u)$ and with a vertex in $\eta(v)$.

Thus, G contains H as a minor if and only if there is an H -expansion in G . (We are mostly concerned with H -expansions when H is a clique, and in this case we also call them *clique-expansions*.) We call the trees $\eta(v)$ ($v \in V(H)$) the *nodes* of the expansion, and denote by $\cup\eta$ the subgraph of G consisting of the union of all the nodes and all the edges $\eta(e)$ ($e \in E(H)$). We say that an H -expansion η is *bipartite* if $\cup\eta$ is bipartite, and *odd* if for every cycle C of $\cup\eta$, the number of edges of C that belong to nodes of the expansion is even. (This unexpected terminology is because it is often said in these circumstances that G contains H as an “odd minor”.) Equivalently, an H -expansion η is odd if we can partition the vertex set of $\cup\eta$ into two subsets A, B such that every edge of each node has one endpoint in A and the other in B , and every edge $\eta(e)$ ($e \in E(H)$) has both endpoints in A or both in B . (We leave checking this equivalence to the reader.)

Hadwiger’s conjecture [6] (see also [7]) states that if a graph contains no K_l -expansion then its chromatic number is at most $l - 1$. This is perhaps the central open problem in graph colouring theory. As shown by Wagner [15], the case $l = 5$ is equivalent to the celebrated four-colour theorem of Appel and Haken [1]. The case $l = 6$ was also shown to follow from the four-colour theorem, by Robertson, Seymour and Thomas [13]. Hadwiger’s conjecture remains open for all larger values of l .

Thomason [14] and Kostochka [8], improving on a result of Mader[10], proved the following:

Theorem 1 *There is a constant c_0 such that for all integers $l \geq 1$, if a graph G has a nonnull subgraph of average degree at least $c_0 l \sqrt{\log l}$ then G contains a K_l -expansion.*

This has a bearing on Hadwiger’s conjecture because of the following well-known observation.

Observation 2 *If $d \geq 0$ is an integer, and every nonnull subgraph of G has average (and hence minimum) degree at most d , then G is $(d + 1)$ -colourable.*

Proof. Let v be a vertex of minimum degree. Inductively there is a $(d + 1)$ -colouring of $G \setminus v$; we extend it to a $(d + 1)$ -colouring of G by choosing a colour for v which appears on none of its neighbours. \square

Thus, the following approximation to Hadwiger’s conjecture is true:

Theorem 3 *If G contains no K_l -expansion then its chromatic number is $O(l \sqrt{\log l})$.* \square

In a similar fashion, we may also bound the chromatic number in terms of the largest bipartite clique-expansion.

Theorem 4 *There is a constant c_1 such that for every integer $l \geq 1$, if G contains no bipartite K_l -expansion, then its chromatic number is at most $c_1 l \sqrt{\log l} + 1$.*

Theorem 4 is an immediate consequence of the following well-known result:

Theorem 5 *If G contains a nonnull subgraph H of average degree at least $2l$ then it contains a nonnull bipartite subgraph of average degree at least l .*

Proof. Choose a two-colouring of H which maximizes the number of bicoloured edges. Let H' be the subgraph of H consisting of all the vertices of H and the bicoloured edges. Clearly, for each vertex v , the degree of v in H' is at least half its degree in H (for if not, then swapping its colour would contradict our choice of bicolouring). \square

Every clique-expansion in a bipartite graph is a bipartite clique-expansion. So, Theorem 1 and Theorem 5 imply:

Theorem 6 *There is a constant $c_1 (= 2c_0)$ such that if G has a nonnull subgraph of average degree at least $c_1 l \sqrt{\log l}$ then it contains a bipartite K_l -expansion.*

Theorem 4 then follows from Observation 2. In this paper we show that an analogous result holds with respect to the largest odd clique-expansion. That is, we show:

Theorem 7 *If G contains no odd K_l -expansion then its chromatic number is $O(l \sqrt{\log l})$.*

This provides evidence for a conjecture of Gerards and Seymour (see [7] page 115) who conjectured that every l -chromatic graph contains an odd K_l -expansion. The key to the proof of Theorem 7 is the following result:

Theorem 8 *If G contains a bipartite K_{16l} -expansion η then either G contains an odd K_l -expansion, or for some set X of vertices with $|X| < 8l$, the (unique) block of $G \setminus X$ that intersects three nodes of η that are disjoint from X is bipartite.*

Note that there is a block U that intersects three nodes of η , because X is disjoint from three of its nodes and there is a cycle in their union. Moreover, U is unique because every pair of nodes of η are joined by an edge. The latter fact also implies that U intersects all of the nodes of η that are disjoint from X . We will prove Theorem 8 in Section 3.

Combining this with Theorem 8 and Theorem 4, we obtain:

Corollary 9 *If G contains no odd K_l -expansion then either its chromatic number is at most*

$$16c_1 l \sqrt{\log 16l} + 1$$

or there exists $X \subseteq V(G)$ with $|X| < 8l$ such that some block of $G \setminus X$ is bipartite and contains more than $8l$ vertices.

It is now an easy matter to deduce Theorem 7 from this corollary. In fact, we have the following strengthening of Theorem 7.

Theorem 10 *Let $c = \lceil 16c_1 l \sqrt{\log 16l} + 24l \rceil$. If G has no odd K_l -expansion then for all $Z \subseteq V(G)$ with $|Z| \leq 16l - 1$, any c -colouring of the subgraph of G induced on Z can be extended to a c -colouring of G .*

Proof. Consider a minimal counterexample (G, Z) to the theorem. If G is $(c - 16l + 1)$ -colourable then we simply colour $G \setminus Z$ with $c - 16l + 1$ colours different from those used on Z . Otherwise, by Corollary 9, there is a set X of vertices of G with $|X| < 8l$, and a bipartite block U of $G \setminus X$ with $|U| > 8l$. For each component K of $G \setminus (X \cup U)$, there is at most one vertex of U that has neighbours in K , since U is a block of $G \setminus X$. Let S_K be the set containing only this vertex if it exists and be empty otherwise, and let $Z'_K = (Z \cap K) \cup X \cup S_K$.

Suppose first that there is no component K of $G \setminus (X \cup U)$ containing at least $8l$ vertices of Z . Then we extend our colouring of Z to a colouring of $Z \cup X$ using $|X|$ colours not used on Z . Since U is bipartite, we can colour $U \setminus Z$ with two of our c colours that are not used on $Z \cup X$. The colouring of $Z \cup X \cup U$ yields a colouring of Z'_K , for each component K of $G \setminus (X \cup U)$. Moreover, $|Z'_K| \leq 16l - 1$, and so by minimality, we can extend the colouring of Z'_K to a c -colouring of the subgraph induced by $K \cup X \cup S_K$. Since there are no edges between K and $G \setminus (X \cup S_K)$, we can combine these colourings (for each K) with our colouring of $Z \cup X \cup U$ to give the desired c -colouring of G .

Finally, suppose that some component K of $G \setminus (X \cup U)$ contains at least $8l$ vertices of Z . Note that $|K \cup Z \cup X \cup S_K| < |V(G)|$, since $|U \setminus S_K| \geq 8l > |Z \setminus K|$. So, by the minimality of G , we can extend the colouring of Z to a c -colouring of the subgraph of G induced by $K \cup Z \cup X \cup S_K$. Let f' be the restriction of this colouring to $X \cup (Z \setminus K) \cup S_K$. Minimality also proves there exists a c -colouring of $G \setminus K$ extending f' . Again, since there are no edges between K and $G \setminus (X \cup S_K)$, combining these two colourings yields the desired c -colouring of G . \square

So it remains to prove Theorem 8. We define a *parity-breaking path* with respect to a bipartite H -expansion η to be a path whose endpoints are in $\cup \eta$ and whose parity differs from the parity of the paths in $\cup \eta$ between them. To transform a bipartite H -expansion into an odd H -expansion we will need many vertex-disjoint parity-breaking paths. Thus, it is not surprising that the following lemma is the crux of the proof of Theorem 8.

Lemma 11 *Let $k \geq 0$ be an integer. For any set S of vertices of a graph G , either*

- (i) *there are k vertex-disjoint paths each of which has an odd number of edges and both its endpoints in S , or*
- (ii) *there is a set X of at most $2k - 2$ vertices such that $G \setminus X$ contains no such path.*

This lemma is of considerable interest in its own right. We prove it in Section 2 and then show that it implies Theorem 8 in Section 3. For a generalization, see [5]. Actually, Theorem 8 has another corollary (Theorem 13 below) which has many important applications. We discuss this result in Section 3; there, by using the proof of Theorem 5 and then applying Theorem 1, we also obtain:

Theorem 12 *There exists a constant c_2 such that every graph containing a K_t -expansion with $t \geq c_2 l \sqrt{\log l}$ also contains a bipartite K_l -expansion.*

Of course, the analogous result with “bipartite” replaced by “odd” cannot hold, because bipartite graphs do not contain an odd K_3 -expansion, since the latter requires an odd cycle, and yet there are bipartite graphs containing K_l -expansions for arbitrarily large l . However, combining Theorem 12 with Theorem 8 yields the following result.

Theorem 13 *There is a constant $c_3 = 16c_2$ such that if G contains a K_t -expansion η where $t = \lceil c_3 l \sqrt{\log 16l} \rceil$ then either G contains an odd K_1 -expansion, or for some set X of vertices with $|X| < 8l$, the (unique) block U of $G \setminus X$ that intersects all the nodes of η disjoint from X is bipartite.*

Finally we discuss some applications of Theorem 13 in Section 4.

2 The Key Lemma

Before proving Lemma 11, we remark that the bound of $2k - 2$ on the size of X is tight. To see this consider the graph formed by a clique C with $2k - 1$ vertices and a stable set S with many more than $2k$ vertices, by adding all possible edges between S and C . Then every odd path with both endpoints in S uses an edge of C , so there do not exist k such paths, vertex-disjoint. On the other hand, every two vertices of S can be combined with every two vertices of C to obtain a path of length three with both endpoints in S . So clearly, a minimum hitting set for this set of odd paths consists of a subset of C of size $|C| - 1$.

We also remark that the proof of Lemma 11 was based on Edmonds' elegant algorithm for determining if there is an odd s - t path in a graph G via testing if an auxiliary graph has a perfect matching.

Proof of Lemma 11. We will construct an auxiliary graph H such that if H has a sufficiently large matching then G has k vertex-disjoint odd paths with their endpoints in S . The Tutte-Berge formula ([2], or [9] Section 3.1) tells us that if a maximum matching in H misses d vertices then there is a set W of vertices of H such that $H \setminus W$ has $|W| + d$ odd components. We will use this structural characterization to find the desired set X if a maximum matching in H is too small to guarantee the existence of k disjoint odd paths with their endpoints in S .

We construct H as follows. For each $v \in V \setminus S$, let v' be a new vertex; let

$$V(H) = V(G) \cup \{v' \mid v \in V \setminus S\},$$

and

$$E(H) = E(G) \cup \{uv' \mid uv \in E(G \setminus S)\} \cup \{vv' \mid v \in V \setminus S\}.$$

In other words, we take the disjoint union of G and a copy of $G \setminus S$, and add an edge between v and its copy, for each $v \in V(G) \setminus S$. Let M be the matching formed by the latter edges. By an M -augmenting path we mean an odd length path in H with end-points in S , and such that its edges alternately belong to $E(H) \setminus M$ and to M . Clearly, the M -augmenting paths in H are in 1-1 correspondence with the odd paths of G that have both their endpoints in S .

By considering the components of $M \cup M'$ for a maximum matching M' of H , we see that if H has a matching of size $|V(G)| - |S| + k$ then G contains k vertex-disjoint odd paths with their endpoints in S .

Thus, it remains to show that if H has no matching of size $|V(G)| - |S| + k$ then there is a set $X \subseteq V(G)$ of size at most $2k - 2$ such that there is no odd path of $G \setminus X$ with both endpoints in S . Therefore suppose that H has no matching of size $|V(G)| - |S| + k$. Since $|V(H)| = 2|V(G)| - |S|$, the Tutte-Berge formula implies that there is a set W of vertices of H such that the number, $oc(W)$, of odd components of $H \setminus W$ is at least $|W| + |S| - 2k + 2$. Choose W maximal with this property; then every component of $H \setminus W$ is odd, as otherwise we could add a vertex in an even component

to W . Let \mathcal{U} be the set of all components of $H \setminus W$. Let Y be the set of vertices of $V(G) \setminus S$ such that both v and v' are in W . Let W' be the set of vertices of $H \setminus W$ joined to W by an edge of M ; so $|W'| = |W| - 2|Y| - |W \cap S|$.

Parity considerations ensure that each $U \in \mathcal{U}$ contains a vertex of $S \cup W'$; let $z(U)$ be some such vertex. Let $Z = \{z(U) \mid U \in \mathcal{U}\}$, and let

$$X' = (S \cup W') \setminus Z.$$

Now $|X'| = |S| + |W| - 2|Y| - |W \cap S| - oc(W)$, and so $|X'| \leq 2k - 2 - 2|Y| - |W \cap S|$. Let

$$X = (X' \cap V(G)) \cup \{v \in V(G) \setminus S \mid v' \in X'\} \cup (W \cap S) \cup Y.$$

Thus, $X \subseteq V(G)$ and $|X| \leq 2k - 2 - |Y| \leq 2k - 2$. We claim that there are no odd paths of $G \setminus X$ with both endpoints in S . It is enough to show that there is no M -augmenting path in $G^* = H \setminus (X \cup \{v' \mid v \in X \setminus S\})$.

Let $v \in V(G) \setminus S$; we say that v' is the *mate* of v and vice versa if one of them belongs to Z (and therefore the other is in W).

(1) $V(G^*) \cap S = S \cap Z$, and $V(G^*) \setminus S$ is the union of all pairs $\{v, v'\}$ with $v \in V(G) \setminus S$ such that either $v, v' \notin W$, or v' is the mate of v . Consequently, every vertex in $V(G^*) \cap W$ has a mate in Z .

The first assertion follows from the definitions of X and G^* . For the second, note that no vertex of $V(G^*) \cap W$ belongs to S , since $W \cap S \subseteq X$, and so the second statement follows from the first.

(2) Z and $V(G^*) \cap W$ are stable sets in G^* .

For Z is stable since its members all belong to different components of $H \setminus W$. But by (1), every vertex in $V(G^*) \cap W$ has a mate in Z . Since $z(U_1)$ and $z(U_2)$ are non-adjacent for all distinct $U_1, U_2 \in \mathcal{U}$, so are the corresponding two vertices of W . This proves (2).

(3) Every M -augmenting path of G^* is contained in $(V(G^*) \cap W) \cup Z$.

Let $U \in \mathcal{U}$. If $z(U) \notin S$, let $C = \{z(U), w\}$, where w is the mate of $z(U)$, and otherwise let $C = \{z(U)\}$. Let $A = U \cap V(G^*) \setminus \{z(U)\}$, and $B = V(G^*) \setminus (A \cup C)$. We claim that C is a cutset separating A and B in G^* . By (1), A is disjoint from S , and so A is paired by edges of M . Let $v \in V(G) \setminus S$ with $v, v' \in U \cap V(G^*) \setminus \{z(U)\}$, and suppose that one of v, v' is adjacent in G^* to some $b \in B$. Since $b \notin A$ and $b \neq z(U)$, it follows that $b \notin U$, and so $b \in W$. By (1), b has a mate in Z , say $z(U^*)$. Since $v, v' \in U$, it follows that $U^* = U$, and so b is the mate of $z(U)$; and therefore $z(U) \notin S$ and $b = w \in C$, contradicting that $b \in B$. Thus there is no such b . This proves our claim that C is a cutset separating A and B in G^* .

Now suppose that P is an M -augmenting path in G^* that contains a vertex of $U \setminus \{z(U)\}$. Since $S \subseteq B \cup C$, it follows that $|C| = 2$ and both vertices in C belong to P ; and so $z(U) \notin S$, and the edge of M containing $z(U)$ is a chord of P . But an M -augmenting path cannot have an edge of M as a chord. Hence every M -augmenting path in G^* is disjoint from $U \setminus \{z(U)\}$. This proves (3).

Combining (2) and (3), we see that every M -augmenting path in G^* is contained in a bipartite graph with bipartition $(V(G^*) \cap W, Z)$. But $V(G^*) \cap W$ is disjoint from S , so there is no M -augmenting (odd) path in G^* . \square

3 The Proof of Theorem 8

We begin with the following observation.

Observation 14 *Let H be a clique of order $2l$, with vertex set $\{a_1, b_1, \dots, a_l, b_l\}$. Suppose that η is a bipartite H -expansion in a graph G , and there are l vertex-disjoint parity-breaking paths P_1, \dots, P_l (with respect to η) such that P_i has one endpoint in $V(\eta(a_i))$ and the other in $V(\eta(b_i))$ and is otherwise disjoint from $\cup \eta$. Then G contains an odd K_l -expansion.*

Proof. Let H' be the clique in H induced on $\{a_1, \dots, a_l\}$. We define an H' -expansion in G as follows. For $1 \leq i \leq l$, let

$$\eta'(a_i) = \eta(a_i) \cup P_i \cup \eta(b_i),$$

and for $1 \leq i < j \leq l$, let

$$\eta'(a_i a_j) = \eta(b_i a_j).$$

Now $\cup \eta$ is bipartite; fix a proper two-colouring of it. For $1 \leq i \leq l$, we convert this to a proper two-colouring of $\eta'(a_i)$ by colouring $\eta(a_i)$ as before and extending this to a proper two-colouring of $\eta'(a_i)$. Our choice of P_i ensures that the vertices of $\eta(b_i)$ have swapped colours. For $1 \leq i < j \leq l$, the edge $\eta'(a_i a_j) = \eta(b_i a_j)$ is bichromatic in the old two-colouring of $\cup \eta$, and is therefore monochromatic in the new colouring. Hence η' is an odd H' -expansion. \square

Given this observation, we can prove the theorem via two applications of Lemma 11. First, however, we need a definition. Let η be an H -expansion in G , and let $v \in V(H)$. A *centre* for $\eta(v)$ is a vertex $t \in V(\eta(v))$ such that for each component T of $\eta(v) \setminus \{t\}$, the number of edges $e \in E(H)$ such that $\eta(e)$ is incident in G with a vertex of T is at most half the number of edges in H incident with v . It is not hard to see that every node $\eta(v)$ has a centre (perhaps more than one). In what follows we assume that for each node, one of its centres has been selected, and we often speak of *the* centre of a node without further explanation.

In particular, the following two lemmas taken together prove Theorem 8.

Lemma 15 *Let η be a bipartite K_{16l} -expansion in G . Then at least one of the following holds:*

1. *there exists $X \subseteq V(G)$ with $|X| < 8l$, such that the block of $G \setminus X$ that intersects all the nodes of η disjoint from X is bipartite;*
2. *there exist $4l$ vertex-disjoint parity-breaking paths with respect to η such that the $8l$ endpoints of these paths are the centres of distinct nodes of η .*

Lemma 16 *Let η be a bipartite K_{12l} -expansion in G . If there exist $4l$ vertex-disjoint parity-breaking paths with respect to η such that the $8l$ endpoints of these paths are the centres of distinct nodes of η , then G contains an odd K_l -expansion.*

Proof of Lemma 15. Let $H = K_{16l}$, and let η be a bipartite H -expansion in G . Let $V(H) = \{h_1, \dots, h_{16l}\}$, and let (A, B) be a bipartition of $\cup\eta$. For $1 \leq i \leq 8l + 2$, we choose s_i to be the centre of $\eta(h_i)$ if this is in A , and to be a new vertex adjacent only to the centre if the centre is in B , thereby constructing an auxiliary graph G^* . We apply Lemma 11 to the set $S = \{s_1, \dots, s_{8l+2}\}$ in G^* , and find either a set of $4l$ vertex-disjoint odd paths with endpoints in S , or a set $X \subseteq V(G^*)$ with $|X| < 8l$ such that there are no odd paths in $G^* \setminus X$ with both endpoints in S . Suppose the first, and let us choose the paths to be minimal; then they have no internal vertices in S . Thus each of these paths, between s_i and s_j say, consists of a parity-breaking path of G between the centres of $\eta(h_i)$ and $\eta(h_j)$ and perhaps a vertex of $V(G^*) \setminus V(G)$ at either end. So in the first case we are done.

We may assume therefore that the second holds; that is, there exists $X \subseteq V(G^*)$ with $|X| < 8l$ such that there are no odd paths in $G^* \setminus X$ with both endpoints in S . Let I be the set of all i with $1 \leq i \leq 16l$ such that X is disjoint from $\{s_i\} \cup V(\eta(h_i))$. Thus $|I| \geq 16l - |X| \geq 8l + 1$, and $|I \cap \{1, 2, \dots, 8l + 2\}| \geq 3$, and we may assume that $1, 2 \in I$. Let U be the unique block of $G \setminus X$ that contains a cycle of $\cup\eta$. Let c_1 be the centre of $\eta(h_1)$, and for each $i \in I$ with $i > 1$, let v_i be the vertex of $\eta(h_1)$ incident in G with $\eta(h_1 h_i)$. Since there are $|I| - 1 \geq 8l$ choices for i , and c_1 is the centre of $\eta(h_1)$, it follows that not all the vertices v_i belong to the same component of $\eta(h_1) \setminus c_1$. Hence there is a path between say v_i, v_j , where $i, j \in I \setminus \{1\}$, containing c_1 . We deduce that $c_1 \in V(U)$, and similarly $V(U)$ contains the centre c_2 of $\eta(h_2)$. There are paths of both parities between any two distinct vertices of a 2-connected non-bipartite graph, since we can link these vertices to an odd cycle by two vertex-disjoint paths. So, if U is not bipartite then there is a parity-breaking path (with respect to η) in G between c_1, c_2 disjoint from X . But possibly adding a vertex of $V(G^*) \setminus V(G)$ at either end of this path in the obvious way yields an odd path of $G^* \setminus X$ with both endpoints in S , a contradiction. Hence U is bipartite. \square

Proof of Lemma 16. Let η be an H -expansion where H is a complete graph with vertex set $\{h_1, \dots, h_{12l}\}$ say, and for $1 \leq i \leq 12l$ let $N_i = V(\eta(h_i))$. Let (A, B) be a bipartition of $\cup\eta$. Let P_1, \dots, P_{4l} be vertex-disjoint parity-breaking paths with respect to η , such that the $8l$ endpoints of these paths are the centres of distinct nodes of η ; and let us choose P_1, \dots, P_{4l} to minimize the number of edges in their union that are not edges of nodes of η .

(1) *Every node of η that contains no endpoint of any of P_1, \dots, P_{4l} is vertex-disjoint from all of P_1, \dots, P_{4l} .*

For suppose that some node does not satisfy this, and let c be its centre. There is a path Q of this node, from c to a vertex x in some P_i , such that $Q \setminus x$ is disjoint from P_1, \dots, P_{4l} . Now there is a parity-breaking path obtained by following Q from c to x and then following P_i from x to one endpoint of P_i . Replacing P_i by this new path contradicts the minimality of our choice. This proves (1).

We relabel so that for $1 \leq i \leq 4l$, P_i has endpoints in N_{2i-1} and N_{2i} and is disjoint from N_j for $j > 8l$. Let J be the subgraph of G induced on

$$V(P_1) \cup \dots \cup V(P_{4l}) \cup N_1 \cup \dots \cup N_{8l}.$$

For $1 \leq i \leq 4l$, let A_i be the set of vertices in $V(J) \cap A$ with a neighbour in $N_{8l+i} \cap B$, and let B_i be the set of vertices in $V(J) \cap B$ with a neighbour in $N_{8l+i} \cap A$. Construct an auxiliary graph

J^* as follows. Add new vertices s_1, \dots, s_{4l} to J , such that for $1 \leq i \leq 4l$, s_i has neighbour set A_i . Then for $1 \leq i \leq 4l$ and each $b \in B_i$, add a new vertex $v(b, i)$ with neighbour set $\{b, s_i\}$, and let $B'_i = \{v(b, i) \mid b \in B_i\}$. Let $S = \{s_1, \dots, s_{4l}\}$.

(2) For every $X \subseteq V(J^*)$ with $|X| < 2l$, there is an odd length path in $J^* \setminus X$ with endpoints in S .

For each vertex of X is in at most one P_i and at most one N_j . So there is at least one value of $i \leq 4l$ such that X is disjoint from $V(P_i) \cup N_{2i-1} \cup N_{2i}$, say $i = 1$. Let c_1, c_2 be the centres of $\eta(h_1)$ and $\eta(h_2)$ respectively. We claim that c_1, c_2 lie in the same block of $J^* \setminus X$. To show this, let I be the set of all i with $3 \leq i \leq 8l$ such that $X \cap N_i = \emptyset$; then $|I| \geq 8l - 2 - |X| \geq 6l - 1$. For each $i \in I$, let u_i be the vertex of N_1 incident in G with $\eta(h_1 h_i)$, and let v_i be the vertex of N_2 incident in G with $\eta(h_2 h_i)$. Also, let $u_0 v_0 = \eta(h_1 h_2)$, where $u_0 \in N_1$ and $v_0 \in N_2$. If there exist distinct $i, j \in I \cup \{0\}$ such that no component of $\eta(h_1) \setminus c_1$ contains both u_i, u_j , and no component of $\eta(h_2) \setminus c_2$ contains both v_i, v_j , then there is a cycle of G with vertex set in the union of N_1, N_2, N_i, N_j containing c_1, c_2 , (where $N_0 = \emptyset$) and since this cycle is disjoint from X it follows that c_1, c_2 belong to the same block of $J^* \setminus X$. Thus we suppose that there are no such i, j . In particular, since $|I| + 1 \geq 2$, there is no $i \in I \cup \{0\}$ such that $u_i = c_1$ and $v_i = c_2$. If some $u_i = c_1$, let T be the component of $\eta(h_2) \setminus c_2$ containing v_i ; then $v_j \in V(T)$ for all $j \in I \cup \{0\}$, contrary to the definition of a centre, since $|I \cup \{0\}| \geq 6l$. Thus $u_i \neq c_1$ and similarly $v_i \neq c_2$ for all $i \in I \cup \{0\}$. Let H be the bipartite graph with vertex set the union of the set of components of $\eta(h_1) \setminus c_1$ and the set of components of $\eta(h_2) \setminus c_2$, with edge set $I \cup \{0\}$, and the natural incidence relation; then H has no matching with cardinality two, and so by König's theorem, some vertex of H is incident with all edges of H . But this is again contrary to the definition of a centre, since $|I \cup \{0\}| \geq 6l$. This proves our claim that c_1, c_2 lie in the same block U of $J^* \setminus X$.

Since there is a parity-breaking path between them which is also disjoint from X , U is non-bipartite. There also must be j, k with $1 \leq j < k \leq 4l$ such that X is disjoint from $\{s_j\} \cup B'_j$ and $\{s_k\} \cup B'_k$. Since there is an edge of G between N_{8l+j} and N_1 , there is an edge of J^* between $\{s_j\} \cup B'_j$ and N_1 ; and similarly there is an edge of J^* between $\{s_k\} \cup B'_k$ and N_2 . Consequently there is a path of $J^* \setminus X$ between s_j and c_1 , consisting of s_j , possibly a vertex of B'_j , and a subpath of $\eta(h_1)$; and similarly there is a path between s_k and c_2 . These two paths are disjoint, and the two centres both belong to U , and since U is not bipartite, it follows that there is an odd length path in $J^* \setminus X$ between s_j and s_k . This proves (2).

By (2) and Lemma 11, we see that J^* contains l vertex-disjoint odd paths with their endpoints in S . By choosing these paths minimal we can ensure that they are internally disjoint from S . By dropping one or two auxiliary vertices at each end of each path, we obtain vertex-disjoint parity-breaking paths of J . Furthermore, if one of the original paths joined say s_j to s_k , then one endpoint x of the corresponding new subpath has a neighbour y in N_j such that xy is bichromatic, and the other endpoint u has a neighbour v in N_k such that uv is bichromatic.

Adding these $2l$ edges yields a set of l vertex-disjoint parity-breaking paths which have their endpoints in distinct elements of $\{N_{8l+1}, \dots, N_{12l}\}$ but which are otherwise disjoint from these nodes. But now applying Observation 14 to the clique-expansion obtained by taking the restriction of η to those nodes which contain endpoints of these paths, we see that G contains an odd clique-expansion of order l , as claimed. \square

Proof of Theorem 12. We remark that for any three graphs A, B, C , if A contains a B -expansion and B contains a C -expansion, then A contains a C -expansion. Similarly, if A contains a bipartite B -expansion, and B contains a C -expansion, then A contains a bipartite C -expansion.

Let c_0 be as in Theorem 1, and let $c_2 = 2c_0$. Let H be a complete graph of order at least $c_2 l \sqrt{\log l}$, and let η be an H -expansion in G . Using some two fixed colours, there are exactly two bicolourings of each node of η . Choose a bicolouring of each node so as to maximize $E(H')$, where H' is the subgraph of H with $V(H') = V(H)$ and $E(H')$ the set of edges $e \in E(H)$ such that the ends of $\eta(e)$ have different colours. Each vertex of H' has degree at least $c_2/2 = c_0$, and so H' contains a K_l -expansion, by 1. But G contains a bipartite H' -expansion, namely the restriction of η to H' ; and so by the remark above, it follows that G contains a bipartite K_l -expansion. \square .

4 Algorithms and Applications

The authors began their study of graphs without large odd clique-expansions, in an attempt to find efficient algorithms for the following decision problems:

1. **k Odd Disjoint Paths:** Given vertices s, t of a graph G , determine whether there are k internally vertex-disjoint odd length s - t paths in G .
2. **k Odd Disjoint Rooted Paths:** Given vertices $\{s_1, \dots, s_k, t_1, \dots, t_k\}$ of a graph G , determine whether there are internally vertex-disjoint odd length paths P_1, \dots, P_k such that P_i has s_i and t_i as endpoints.
3. **k Odd Disjoint Cycles:** Determine whether there are k vertex-disjoint odd cycles in a graph G .
4. **Odd H -Expansion:** Determine whether G contains an odd H -expansion.

We remark that the variants of these problems in which we drop the parity condition are all solvable in polynomial time (with the parameters k, H fixed). The k Disjoint Paths problem can be solved efficiently even if k is part of the input. In contrast k Disjoint Rooted Paths is NP-complete unless k is fixed, in which case it can be solved in polynomial time. The algorithm required the development of a complicated structure theorem characterizing graphs without large clique minors (see [12]). H -Minor Containment is also NP-complete if H is part of the input, as the Hamilton Cycle problem is a special case. However, the techniques of Robertson and Seymour can be used to solve this problem efficiently for fixed H . The k Disjoint Cycles problem is also NP-complete as it contains the Vertex Cover by Triangles problem as a special case. However, it is solvable in linear time for fixed k because the Erdős-Posa property holds for cycles and it follows that graphs without k vertex-disjoint cycles have tree-width bounded by a function of k (see [4] for details and extensions).

The odd variants of the last three of these problems are also NP-complete if k or H is part of the input. Hamilton Cycle is the special case of Odd H -Expansion where H is a cycle with the same number of vertices as G ; Vertex Cover by Triangles is the special case of k Odd Disjoint Cycles where $n = 3k$. Finally, it is easy to transform an instance of k -Disjoint Rooted Paths to an instance of

k Odd Disjoint Rooted Paths. We simply add, for each edge, a path of length two connecting its endpoints. We do not know if k Odd Disjoint Paths is NP-complete if k is part of the input.

We have an efficient algorithm to solve the k Disjoint Odd Cycles problem for fixed k , which we are currently writing up. We believe that we can efficiently solve the other three problems. However, in doing so we have to rework much of the structure theory for graphs with no large clique minor. This is a daunting but doable task, which we have begun.

To this end, we note that our proof of Lemma 11 is algorithmic. That is, it provides an algorithm for the following problem (where k is part of the input and need not be fixed):

5. **Odd Disjoint S -Paths:** Given a set S of vertices in a graph G and an integer k find either k vertex-disjoint odd paths with their endpoints in S or a set X of at most $2k - 2$ vertices which meet all such paths.

In our proof, we built an auxiliary graph and determined whether or not the desired paths exist by solving the maximum matching problem on this auxiliary graph. The rest of the proof is clearly algorithmic. In order to choose X so that all its components are odd, we may need to add $O(|V(G)|)$ vertices to X , and after each addition it takes $O(|E(G)|)$ time to recompute the components. The rest of the proof can easily be implemented in $O(|E(G)|)$ time. So, the total time complexity is $O(|V(G)||E(G)| + MM)$ where MM is the time taken to solve maximum matching on a graph with at most $2|V(G)|$ vertices and $2|E(G)| + |V(G)|$ edges. Actually we can do slightly better than this because we can start with the matching M , so if our maximum matching algorithm is an iterative augmenting algorithm than we need only perform k augmenting steps.

The proof of Theorem 8 can also be made constructive, that is we can provide an efficient algorithm which given a bipartite K_{16l} -expansion, either finds a odd K_l -expansion or the set X with the properties specified in the statement of the theorem. A glance at the proof shows that we essentially apply Lemma 11 twice to two auxiliary graphs as well as doing some straightforward cleanup operations so this algorithm has the same time complexity as our algorithm for Odd Disjoint S -Paths.

Finally, there is a polynomial-time algorithm to obtain the colouring extensions guaranteed by Theorem 10. We first repeatedly remove a minimum degree vertex from $G \setminus Z$ as long as its degree is less than $c - 16l$. If we remove the entire graph then $G \setminus Z$ is greedily $c - 16l$ colourable so we colour it using colours which do not appear on Z . Otherwise, we find a subgraph of G with minimum degree at least $c - 16l$. We find a bipartite subgraph within this subgraph of average degree at least $c/2 - 8l$ (this is easy to do greedily; we colour the vertices in arbitrary order, always colouring a vertex so as to maximize the number of bichromatic edges out of it and into some already coloured vertex). Next, we find a (bipartite) K_{16l} -expansion within this bipartite graph using an algorithmic version of Theorem 1. Now, we apply our algorithmic version of Theorem 8 to find a set X as in the statement of that theorem (we assume our input has no odd K_l -expansion so this is the only possible output).

At this point in the proof of Theorem 10, we show that solutions to some subproblems can be combined to obtain solutions to our original problem. So our algorithm will recurse at this point, solving the specific subproblems. We note that the sum over all subproblems (G', Z') of $|V(G')| - |Z'|$ is at most $|V| - |Z|$ because no vertex is freely colourable in two distinct subproblems. It follows trivially by induction that we perform at most $2(|V| - |Z|) - 1$ iterations when solving an instance

provided $|V| > |Z|$. Thus, we perform a polynomial number of iterations each of which takes polynomial time and so the algorithm is polynomial. We make no attempt to optimize it.

References

- [1] Appel K. and Haken W. (1977), “Every planar map is 4-colourable, Part I: Discharging”, *Illinois J. Math.* 21:429-490.
- [2] Berge C. (1958), “Sur le couplage maximum d’un graphe”, *Comptes Rendus de l’Academie de Sciences de Paris, Series 1 Mathematique*, 247:258-259.
- [3] Dejter I., and Neumann-Lara V. (1985), “Unboundedness for generalized odd cycle traversability and a Gallai conjecture”, paper presented at the *Fourth Caribbean Conference on Computing*, Puerto Rico.
- [4] Erdős P., and Pósa L. (1965), “On independent circuits contained in a graph”, *Canadian Journal of Mathematics*, 17:347-352.
- [5] Chudnovsky M., Geelen J., Goddyn L., Lohman M., and Seymour P. (2004), “Packing non-zero A-paths in group-labeled graphs”, preprint.
- [6] Hadwiger H. (1943), “Über eine Klassifikation der Streckenkomplexe”, *Vierteljahrsschrift der naturforschenden Gesellschaft in Zurich*, 88:133-142.
- [7] Jensen T., and Toft B. (1995), *Graph Colouring Problems*, Wiley, Chichester UK.
- [8] Kostochka A. (1984), “Bounds on the Hadwiger number of graphs by their average degree”, *Combinatorica*, 4:307-316.
- [9] Lovász L., and Plummer M. (1986), *Matching Theory*, North Holland, Amsterdam.
- [10] Mader W. (1968), “Homomorphiesätze für Graphen”, *Math Ann.*, 178:154-168.
- [11] Reed B. (1999), “Mangoes and blueberries”, *Combinatorica*, 19:267-296.
- [12] Robertson N., and Seymour P. (1995), “Graph Minors. XIII. The disjoint paths problem”, *J. Combinatorial Theory Ser. B*, 63:65-110.
- [13] Robertson N., Seymour P., and Thomas R. (1993), “Hadwiger’s conjecture for K_6 -free graphs”, *Combinatorica*, 13:279-361.
- [14] Thomason A. (1984), “An extremal function for contractions of graphs”, *Math. Proc. Camb. Phil. Soc.*, 95:261-265.
- [15] Wagner K. (1937), “Über eine Eigenschaft der ebenen Komplexe”, *Math Ann.*, 114:570-590.