Reconfiguration graphs for vertex colourings of chordal and chordal bipartite graphs *†

Marthe Bonamy[1]‡, Matthew Johnson[2], Ioannis Lignos[2], Viresh Patel[2], and Daniël Paulusma[2]

[1] Laboratoire d'Informatique, de Robotique et de Microélectronique de Montpellier, France, marthe.bonamy@lirmm.fr

[2] School of Engineering and Computing Sciences, Durham University, United Kingdom, {matthew.johnson2,i.m.lignos,viresh.patel,daniel.paulusma}@durham.ac.uk

Abstract

A k-colouring of a graph $G = (V, E)$ is a mapping $c : V \to \{1, 2, \ldots, k\}$ such that $c(u) \neq c(v)$ whenever uv is an edge. The reconfiguration graph of the k-colourings of G contains as its vertex set the k-colourings of G, and two colourings are joined by an edge if they differ in colour on just one vertex of G. We introduce a class of k-colourable graphs, which we call *k-colour-dense* graphs. We show that for each k-colour-dense graph G, the reconfiguration graph of the ℓ-colourings of G is connected and has diameter $O(|V|^2)$, for all $\ell \geq k + 1$. We show that this graph class contains the k-colourable chordal graphs and that it contains all chordal bipartite graphs when $k = 2$. Moreover, we prove that for each $k \geq 2$ there is a k-colourable chordal graph G whose reconfiguration graph of the $(k + 1)$-colourings has diameter $\Theta(|V|^2)$.

Keywords reconfigurations, graph colouring, graph diameter, chordal graphs.

1 Introduction

The reconfiguration graph of the k-colourings of a graph G contains as its vertex set the k-colourings of G, and two colourings are joined by an edge in the reconfiguration graph if they differ in colour on just one vertex of G. In this paper, we determine sufficient conditions for the reconfiguration graph

*This paper is supported by EPSRC Grants EP/E048374/1 and EP/F064551/1.

†A preliminary version of this paper has been presented at EuroComb 2011.

‡Supported by École Normale Supérieure de Lyon

to have a diameter that is at most quadratic in the number of vertices. We give examples of graph classes, such as chordal graphs and chordal bipartite graphs, that satisfy these conditions and describe a class of graphs that show that our quadratic bound is tight.

This work continues the study of *reconfiguration problems* found in a number of recent papers. We can define the reconfiguration graph for any search problem: the vertex set contains all solutions to the problem; the edge set is defined by a symmetric adjacency relation on the solutions which is normally chosen to represent a smallest possible change in the solution. To date, the study of reconfiguration graphs has focussed on the computational complexity of the problems of deciding whether the reconfiguration graph is connected, and deciding whether it contains a path between two given solutions. Problems studied include boolean satisfiability [10], graph colouring [3, 5, 6], shortest path [2, 14], and independent set, clique and others [13].

Reconfiguration problems have diverse motivations. First, they represent an application in which it is necessary to move between solutions passing step-by-step through only feasible solutions (such as when the solution represents an allocation by a supplier to customers as in the Power Supply problem [13]). Second, they can represent the evolution of a genotype where only single mutations can occur and all genotypes must be above a certain fitness threshold. Finally, an understanding of the geometry of the solution space can provide insight into the performance of algorithms and heuristics [1].

A fundamental problem is to characterise the relationship between the complexity of reconfiguration problems and search problems. Considering the problem of finding paths between solutions, previous results often follow a certain pattern: problems in P beget reconfiguration problems that are also in P; NP-complete problems have PSPACE-complete reconfiguration problems. There are exceptions such as the shortest path reconfiguration problem being PSPACE-hard [2].

Also of interest is finding *shortest* paths between solutions. The diameter of the reconfiguration graph provides an upper bound. This is also related to the complexity of finding paths in the reconfiguration graph between given solutions since paths of polynomial length in the reconfiguration graph are certificates for the problem being in NP. For any graph, the diameter of the reconfiguration graph of its 3-colourings has been shown to be at most quadratic (in the number of vertices of the input graph) if the reconfiguration graph is connected [6]. Although there are cases where the reconfiguration graph is not connected but contains components of superpolynomial diameter [3], there is no known example of a family of graphs for which the reconfiguration graph of k-colourings is connected but does not have (at most) quadratic diameter.

A good place to start when thinking about the above question is to consider graphs of bounded degeneracy. It is well known that graphs of degeneracy k are $(k+1)$-colourable. Bonsma and Cereceda [3] showed that if G is a graph of degeneracy k, then R_G^{k+2}, the reconfiguration graph of $(k+2)$-colourings of G, is connected. In light of the previous paragraph, we are naturally led to ask whether R_G^{k+2} has quadratic diameter; indeed it is conjectured [3] that R_G^{k+2} has cubic diameter, although this is modified to quadratic [4]. Our work includes an important class of k-degenerate graphs, namely $(k+1)$-colourable chordal graphs, for which we show the conjecture to be true.

2 Preliminaries

In this section we give some basic terminology and notation; for any undefined terminology in the paper we refer the reader to the textbook of Diestel [7].

We consider undirected finite graphs that have no loops and no multiple edges. A graph is denoted $G = (V, E)$, where V is the set of vertices and E is the set of edges. For a subset $S \subseteq V$, the graph $G[S]$ denotes the subgraph of G *induced by* S, i.e., the graph with vertex set S and edge set $\{uv \in E \mid u, v \in S\}$. We write $G - S = G[V \setminus S]$. The *disjoint union* of two vertex-disjoint graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, which we denote by $G_1 \cup G_2$, is the graph with vertex set $V_1 \cup V_2$ and edge set $E_1 \cup E_2$.

The set of *neighbours* of a vertex u in a graph G is denoted $N_G(u) = \{v \mid uv \in E\}$. If u has no neighbours, then we say that u is an *isolated* vertex. If u and v are adjacent and have no other neighbours, then the edge uv is called an *isolated* edge.

A *(vertex) colouring* of a graph $G = (V, E)$ is a mapping $c : V \to \{1, 2, \ldots\}$ such that $c(u) \neq c(v)$ whenever $uv \in E$. Here, $c(u)$ is referred to as the *colour* of u. We write $c(U) = \{c(u) \mid u \in U\}$ for $U \subseteq V$. Then a *k-colouring* of G is a colouring c of G with $c(V) \subseteq \{1, \ldots, k\}$. If G has a k-colouring, then G is called *k-colourable.* The *chromatic number* of G denoted χ_G is the smallest value of k for which G is k-colourable. If G is 2-colourable, then G is also called *bipartite.* We denote the reconfiguration graph of the k-colourings of a graph G by R_G^k. Recall that R_G^k contains as its vertex set the k-colourings of G, and two colourings are joined by an edge in R_G^k if they differ in colour on just one vertex of G.

The n-vertex *path* is the graph with vertices $v_1, \ldots, v_n$ and edges $v_i v_{i+1}$ for $i = 1, \ldots, n-1$. If $v_n v_1$ is also an edge, then we obtain the n-vertex *cycle*. The *length* of a path or a cycle is the number of its edges. A graph is called *connected* if, for every pair of distinct vertices v and w, there exists a path connecting v and w. A maximal connected subgraph D of a graph is called a *connected component* (or just *component*) of G; we shall often abuse notation by denoting both the connected component and its vertex set by D. A *separator* of a graph $G = (V, E)$ is a set $S \subset V$ such that $G - S$ has more connected components than G; if two vertices u and v that belong to the same connected component in G are in two different connected components of $G - S$, then we say that S separates u and v. We say that we *identify* two vertices u and v if we replace them by a new vertex adjacent to all neighbours of u and v.

A *tree* is a connected graph with no cycles. A *clique* is a graph where every pair of vertices is joined by an edge. The size of a largest clique in G is denoted ω_G. A *perfect* graph is a graph in which $\chi_{G'} = \omega_{G'}$ for every (not necessarily proper) induced subgraph of G.

3 Sufficient Conditions for Quadratic Diameter

In this section, we introduce the class of *k-colour-dense* graphs, and we show by induction in Theorem 2 that, for every k-colour-dense graph G, the diameter of R_G^ℓ is quadratic in the size of G for all $\ell \geq k+1$. Indeed, the definition of k-colour-dense graphs is recursive and has been formulated in order to facilitate our inductive method. For this reason, it is difficult to establish precisely which graphs are k-colour-dense; however, in the next section, we shall show that, for example, k-colourable chordal graphs are k-colour-dense.

For a fixed positive integer k, we say that a k-colourable graph G on n vertices is *k-colour-dense* if either

(i) G is the disjoint union of cliques, each of which has at most k vertices, or

(ii) G has a separator S, and $G - S$ has components D and D' with vertices $u \in D$ and $v \in D'$ such that

  (a) $|D| = 1$ or $|D \cup S| \leq k$,

  (b) $S \subseteq N(v)$, and

(c) identifying u and v in G results in a k-colour-dense graph G'.

We make the following easy observation for use in Section 4.1

Proposition 1. *If G_1 and G_2 are k-colour-dense graphs, then $G_1 \cup G_2$, the disjoint union of G_1 and G_2, is k-colour-dense.*

Proof. The proof is by induction on the total number of vertices in G_1 and G_2. If G_1 and G_2 are both the disjoint union of cliques, then the claim holds trivially, so assume that G_1 is not the disjoint union of cliques. Thus G_1 has a separator S, components D_1 and D_2, and vertices u and v as in part (ii) of the definition of k-colour-dense graphs; in particular, G'_1, the graph obtained from G_1 by identifying u and v, is k-colour dense. Thus, by induction, the disjoint union of $G'_1 \cup G_2$ is k-colour dense. Thus S, D, D', u, v also fulfills part (ii)(c) of the definition of k-colour-dense graphs when applied to $G_1 \cup G_2$ (and they obviously still satisfy (ii)(a) and (ii)(b)). □

We define the *ℓ-colour diameter* of a graph G to be the diameter of R^ℓ_G.

Theorem 2. *For an integer $k \geq 1$, let G be a k-colour-dense graph on n vertices. Then, for all $\ell \geq k+1$, the ℓ-colour diameter of G is at most $2n^2$.*

Note that $\ell \geq k+1$ is necessary in the above theorem because, for example, the reconfiguration graph of the k-colourings of a clique on k vertices consists of $k!$ isolated vertices.

Proof. Let $k \geq 1$ be an integer and let G be a k-colour-dense graph on n vertices. We assume $\ell = k+1$; the proof for $\ell > k+1$ is similar. We prove the following claim which immediately implies the theorem.

Claim 1. Let α and β be two $(k+1)$-colourings of G. Then we can transform α to β by recolouring every vertex of G at most $2n$ times.

There are two cases to consider corresponding to the two conditions in the definition of k-colour dense graphs.

We first suppose that G is a disjoint union of cliques and describe how to recolour from α to β. We recolour the disjoint cliques one at a time. Given a clique of G with vertices $v_1, \ldots, v_r$, $(r < k+1)$, we consider the vertices in order; once we have $v_1, \ldots, v_{i-1}$ coloured with colours $\beta(v_1), \ldots, \beta(v_{i-1})$ respectively, we try to recolour v_i with $\beta(v_i)$. We are only prevented from doing this directly if there is a vertex v_j with $j > i$ that is presently coloured $\beta(v_i)$. In this case we first recolour v_j with an unused colour (such a colour exists since $r < k+1$) and then colour v_i with $\beta(v_i)$. When the whole clique is coloured with β each v_j has been recoloured at most $j \leq r \leq 2n$ times.

We now consider the case where G is not a disjoint union of cliques but satisfies condition (ii) of the definition of k-colour dense. We use induction on the number of vertices. Let S, D, D', $u \in D$ and $v \in D'$ be as in condition (ii). We first show how to transform α into some $(k+1)$-colouring α' satisfying $\alpha'(u) = \alpha'(v)$, by recolouring each vertex of G at most once. Suppose that $\alpha(u) \neq \alpha(v)$. If we can immediately recolour u with $\alpha(v)$, then we do this to obtain the desired colouring α'. If not, then

$$W := \{w \in N_G(u) \mid \alpha(w) = \alpha(v)\} \subseteq N_G(u)$$

must be non-empty. Since $u \in D$, and D is a component of $G - S$, we have $W \subseteq N_G(u) \subseteq D \cup S$. However, every vertex of W is coloured $\alpha(v)$ and no vertex of S is coloured $\alpha(v)$ (since every vertex of S is adjacent to v by condition (ii)(b)), so $W \subseteq D$. Now, for each $w \in W \subseteq D$, we have $N_G(w) \subseteq D \cup S$; thus $|N_G(w)| \leq |D \cup S| \leq k$ by condition (ii)(a) (note that $|D| \neq 1$ since D contains u and the non-empty set $W \subseteq N(u)$). Hence, each vertex of W can be successively recoloured with some colour not used in its neighbourhood. After this we recolour u with $\alpha(v)$ and we do not recolour any other vertices of G. Thus we have recoloured each vertex of G at most once and transformed α to a new $(k+1)$-colouring α' where $\alpha'(u) = \alpha'(v)$. By the same argument, we can transform β to a $(k+1)$-colouring β' with $\beta'(u) = \beta'(v)$. Changing α to α' and β to β' together require that each vertex of G is recoloured at most twice.

We now identify u and v. This leads to a new vertex u' and a graph G' that is k-colour-dense by condition (ii)(c). We can consider α' and β' to be colourings of G' by defining $\alpha'(u') = \alpha'(u) = \alpha'(v)$ and $\beta'(u') = \beta'(u) = \beta'(v)$, respectively. We can transform α' into β' on G' using at most $2(n-1)$ recolourings for each vertex (by application of either the induction hypothesis or the previous case depending on whether G' satisfies the first or second condition of the definition of k-colour dense). Thus we can transform α' into β' on G by simulating each recolouring of u' by a recolouring of u and v in G, i.e., every time we recolour u' in G' we apply the same recolouring to u and then immediately to v in G. Thus transforming α' to β' in G requires that each vertex of G is recoloured at most $2(n-1)$ times, and transforming α to α' and β' to β requires at most two additional recolourings of each vertex, resulting in a total of at most $2(n-1)+2 = 2n$ recolourings of each vertex, as required. This completes the proof of the claim and of Theorem 2. $\square$

4 Graph Classes

We show that k-colourable chordal graphs are k-colour-dense for every fixed integer $k \geq 1$ and that chordal bipartite graphs are 2-colour-dense. Hence, these graphs satisfy the necessary condition in Theorem 2 and consequently have a quadratic ℓ-colour diameter for $\ell \geq k+1$ and $\ell = 3$, respectively.

4.1 Chordal graphs

A *chordal* graph is a graph with no induced cycle of length more than 3. Let $G = (V, E)$ be a graph, let $\mathcal{K}$ be the set of maximal cliques of G, and for $v \in V$, let $\mathcal{K}_v$ be the set of maximal cliques of G containing v. A *clique tree* $\mathcal{T}$ of a (connected) graph G is a tree whose vertex set is $\mathcal{K}$ and whose edges are such that $\mathcal{T}[\mathcal{K}_v]$ is connected (i.e. forms a subtree) for all $v \in V$. In this context, the maximal cliques of G are also called *bags* of $\mathcal{T}$.

The next lemma is well known.

Lemma 3 ([12]). *A connected graph is chordal if and only if it has a (not necessarily unique) clique tree.*

The next lemma is also well known (see e.g. [9]).

Lemma 4. *If G is a chordal graph then $\omega_G = \chi_G$.*

Next we prove some properties of chordal graphs and clique trees that we shall require. The first property is well known [8], and the second one has probably been used before, but we give proofs for completeness.

Lemma 5. *Let G be a connected chordal graph that has a clique tree $\mathcal{T}$, where $\mathcal{T}$ has at least two vertices. Let K be a leaf of $\mathcal{T}$ and let K' be the unique neighbour of K in $\mathcal{T}$. We have the following properties.*

(i) We have that $S := K \cap K'$ is a separator of G, and $D := K \backslash S$ is non-empty and a connected component of $G - S$.

(ii) There exists $u \in K \backslash K' = K \backslash S = D$ and $v \in K' \backslash K$ such that, if G' is obtained from G by identifying u and v, then G' is chordal and $\omega_{G'} \leq \omega_G$ (so $\chi_{G'} \leq \chi_G$ by Lemma 4).

We remark that the above lemma holds more generally even if K is not a leaf of $\mathcal{T}$, but the proof in our case is slightly simpler.

Proof. (i) Fix any $u \in D := K \backslash S = K \backslash K'$; such a vertex exists since otherwise $K \subseteq K'$, contradicting the maximality of K. Fix any $z \in G - K$. Let P be a path of G from u to z with vertices $u = a_0, a_1, \ldots, a_r, a_{r+1} = z$ in order. Let $a_i a_{i+1}$ be the first edge of P not in K. Then $a_i a_{i+1}$ is an edge of some maximal clique $K^* \neq K$. Furthermore $a_i \in K$ since either $a_i = u$ or $a_{i-1} a_i$ is an edge of K. We deduce that $K, K^* \in \mathcal{K}_{a_i}$. Since $\mathcal{T}[\mathcal{K}_{a_i}]$ is connected and the only neighbour of K is K', we have $K' \in \mathcal{K}_{a_i}$. Thus $a_i \in K \cap K' = S$ and so P passes through S. So every path from $u \in K \backslash S$ to any vertex $z \notin K$ passes through S. Hence S is a separator of G, and $K \backslash S =: D$ (which is a clique) is a connected component of $G - S$.

(ii) Fix any $u \in K \backslash S = K \backslash K'$ and $v \in K' \backslash K$; such vertices exist by the maximality of K and K'. Let G' be the graph obtained by identifying u and v, and let u' be the new vertex of G' that results. Suppose for a contradiction that G' is not chordal. Then G' has an induced k-cycle for some $k \geq 4$; this cycle necessarily contains u' since otherwise G would contain an induced k-cycle. Therefore in G' there is a path with vertices $u, b_1, \ldots, b_{k-1}, v$ (in order) such that identifying u and v gives an induced cycle. Thus the path can have no chords except possibly ub_{k-1} or $b_1 v$. However both of those chords would give an induced k-cycle in G, so we can assume that P is an induced path (of length $k \geq 4$). But, since S separates u and v (by part (i) of the lemma), P must pass through S, and since every vertex of $S = K \cap K'$ is adjacent to both u and v, P cannot be an *induced* path.

Finally, suppose for a contradiction that G' has a $(k+1)$-clique. The clique necessarily contains u'; otherwise it would also be a $(k+1)$-clique of G. Thus in G, there is a k-clique L such that $L \subseteq N(u) \cup N(v)$. Fix vertices $a \in L \backslash N(u)$ and $b \in L \backslash N(v)$ (a, b exist because otherwise we have a $(k+1)$-clique of G). We know that $S \subseteq N(u) \cup N(v)$, so that $a, b \notin S$. We also know S separates u and v, and yet u, a, b, v is a path from u to v in $G - S$, a contradiction. □

We use Lemma 5 in the proof of the following result.

Theorem 6. *For each fixed integer $k \geq 1$, every k-colourable chordal graph is k-colour-dense.*

Proof. Let $G = (V, E)$ be a k-colourable chordal graph on n vertices. We show by induction on n that G is k-colour-dense. We may assume that G is connected since otherwise, each component of G is k-colour-dense (by induction), and so G is k-colour dense by Proposition 1. We may also assume that G is not a clique, since then it is trivially k-colour-dense.

By Lemma 3, G has a clique tree $\mathcal{T}$. Since G is not a clique, G has at least two maximal cliques, so $\mathcal{T}$ has at least two vertices. Let K be a leaf of $\mathcal{T}$, and let K' be the unique neighbour of K. By Lemma 5, $S := K \cap K'$ is a separator of G, $D := K \backslash S$ is a connected component of $G - S$, and there exist two vertices $u \in D$ and $v \in K' \backslash K \subseteq V \backslash (D \cup S)$ such that identifying u and v gives a graph G' that is chordal and $\chi_{G'} \leq \chi_G \leq k$. Set D' to be the connected component of $G - S$ containing v.

Now, for G, it is easy to check that S, D, D', u, v satisfy conditions (ii) in the definition of k-colour-dense graphs. Condition (ii)(a) is satisfied because $D \cup S = K$ and so $|D \cup S| \leq |K| \leq k$. Condition (ii)(b) is satisfied because $v \in K'$ and $S \subseteq K'$, so that $S \subseteq N(v)$. Condition (ii)(c) is satisfied because identifying u and v in G gives a k-colourable chordal graph G', which is k-colour-dense by the induction hypothesis. □

4.2 Chordal bipartite graphs

A *chordal bipartite* graph is a bipartite graph with no induced cycle of length more than 4. It is a misnomer since chordal bipartite graphs are only chordal if they are trees. We show that chordal bipartite graphs are 3-colour-dense by proving that a more general class of graphs is 3-colour-dense. Let us call a graph *semi-false* if it can be constructed from a set of one or more isolated vertices by a sequence of the following two operations, namely adding a pendant vertex and adding a semi-false twin. Here, a *pendant* vertex in a graph is a vertex of degree one, and a vertex u is a *semi-false* twin of another vertex v if $N(u) \subseteq N(v)$. Note that adding a pendant vertex u is a special case of adding a semi-false twin, unless u is added as the neighbour of an isolated vertex.

In order to show that every chordal bipartite graph is semi-false we need the following terminology. A vertex u in a bipartite graph G is *weakly simplicial* if its neighbours can be labelled $v_1, \ldots, v_t$ such that $N(v_i) \subseteq N(v_{i+1})$ for $i = 1, \ldots, t-1$. Uehara [16] showed the following, which also follows from results of Hammer, Maffray, and Preismann [11]; see Pelsmajer, Tokazy, and West [15].

Lemma 7 ([11, 16])**.** *A bipartite graph is chordal bipartite if and only if every induced subgraph has a weakly simplicial vertex.*

We use Lemma 7 in the proof of the following theorem.

Theorem 8. *The class of semi-false graphs is a proper superclass of the class of chordal bipartite graphs.*

Proof. We first give an example of a semi-false graph G^* that is not chordal bipartite. Start with a vertex u_1 and add three pendant vertices u_2, u_3, u_4, each with (unique) neighbour u_1. Then add two semi-false twins u_5 and u_6 of u_1 with neighbours u_2, u_3 and u_3, u_4, respectively. Finally add a semi-false twin u_7 of u_3 with neighbours u_5 and u_6. Because $u_1, u_2, u_4, u_5, u_6, u_7$ induce a 6-vertex cycle in G^*, we find that G^* is not chordal bipartite.

We now show by induction on n that every chordal bipartite graph G on n vertices is semi-false. The case $n = 1$ is trivial. Let $n \geq 2$, let G be a chordal bipartite graph on n vertices, and assume that every chordal bipartite graph with $n - 1$ vertices is semi-false. If we can show that G can be obtained from a semi-false graph G' by adding a pendant vertex or a semi-false twin the theorem will follow. Note that any graph obtained from G by removing a vertex is chordal bipartite and so, by the induction hypothesis, semi-false.

As a graph containing only isolated vertices is semi-false, we assume that G has a component D containing at least 2 vertices. Lemma 7 tells us that D has a weakly simplicial vertex u, the neighbours of which can be labelled $v_1, \ldots, v_t$, $t \geq 1$, such that $N(v_i) \subseteq N(v_{i+1})$ for $i = 1, \ldots, t-1$.

First suppose that $t = 1$. Then let $G' = G - u$. Thus G is obtained from G' by adding u as a pendant vertex.

Now suppose that $t \geq 2$. Then let $G' = G - v_1$. Therefore G is obtained from G' by adding v_1 as a semi-false twin of v_2. $\square$

We note that the class of semi-false graphs does not contain the class of chordal graphs; this can be seen by taking any clique on 3 or more vertices.

We now show that semi-false graphs are bipartite.

Proposition 9. *Every semi-false graph G is 2-colourable.*

Proof. If G contains only isolated vertices the proposition is true. Otherwise G can be obtained from a graph G' by adding a vertex u that is either pendant or a semi-false twin. Using induction, we can assume that G' has a 2-colouring. We show how to extend it to G by colouring u. If u is pendant, we colour it with the colour that is not used on its unique neighbour. If u is a semi-false twin, then all its neighbours have a common neighbour v. We can therefore colour u with the colour used on v. $\square$

We conclude this section by showing that every semi-false graph G is 2-colour-dense.

Theorem 10. *Every semi-false graph is 2-colour-dense.*

Proof. We prove by induction on n that if $G = (V, E)$ is a semi-false graph on n vertices then it is 2-colour-dense. The claim is trivially true if $n = 1$.

If G is a semi-false graph on n vertices, then we know by Proposition 9 that G is 2-colourable. Recall that G is constructed from a set U of isolated vertices by a sequence of pendant-vertex and semi-false-twin operations. Let u be the last vertex added to G either as a pendant vertex or a semi-false twin (if there is no such vertex, then we have $G = (U, \emptyset)$, which is trivially 2-colour dense). If u is a pendant vertex, we may assume that u is an end vertex of an isolated edge $e = uu'$ of G (since otherwise we can consider u to be a semi-false twin of another vertex). Then $G[\{u, u'\}] = (\{u, u'\}, \{e\})$ is 2-colour-dense, $G[V \setminus \{u, u'\}]$ is 2-colour-dense by induction, so G is 2-colour dense by Proposition 1.

Thus we may assume u is a semi-false twin of some other vertex v of G. We take $S = N(u)$, $D = \{u\}$ and we let D' be the component of $G - S$ containing v. Then S is a separator of G (separating u from v) and $|D| = 1$; hence, condition (ii)(a) in the definition of 2-colour-dense is satisfied. Because $S = N(u) \subseteq N(v)$, condition (ii)(b) is satisfied. Finally, identifying u and v in G to form G' is equivalent to deleting u from G. Thus G' is a semi-false graph (obtained from U by performing the same operations as used for G, except the last). Since G' is 2-colour-dense (by induction) we see that condition (ii)(c) is satisfied. This completes the proof of Theorem 10. $\square$

5 Lower Bounds

We prove that the bound in Theorem 2 is asymptotically tight up to a constant factor for every k. To be more precise, for $k = 2$, we show that the 3-colour diameter of a path on n vertices is $\Theta(n^2)$. (Note that a path is chordal bipartite, and as such it is 2-colour-dense due to Theorems 8 and 10.) Apart from one subtlety, our result employs very similar techniques to [6], where it is shown that a path on n vertices with an appended triangle has two 3-colourings with quadratic separation. Note however that this example has a disconnected reconfiguration graph and hence infinite diameter.

For each fixed $k \geq 3$ and every $n \geq k$, we give an example of an n-vertex, k-colour-dense graph $G_k(n)$ with $(k+1)$-colour diameter $\Theta(n^2)$. We believe that these are the first examples of graphs with quadratic $k+1$-colour diameter. These examples are easily derived from the path.

Theorem 11. *The* 3*-colour diameter of a path on* n *vertices is* $\Theta(n^2)$*.*

Proof. We have already seen that the 3-colour diameter of a path on n vertices is at most $2n^2$ by Theorem 2 and recalled that a path is 2-colour-dense. It remains only to show a quadratic lower bound.

Let P be a path on n vertices $v_1, \dots, v_n$ for some integer $n \geq 2$. Let the $n-1$ edges of P be $e_1, \dots, e_{n-1}$, where $e_i = v_i v_{i+1}$ for $i = 1, \dots, n-1$. We define edge weights $w(e_i) = \min(i, n-i)$ for $i = 1, \dots, n-1$. Given a 3-colouring c of P and an edge $e_i = v_i v_{i+1}$, we define

$$z_c(e_i) = \begin{cases} 1 & \text{if } (c(v_i), c(v_{i+1})) = (1,2), (2,3), \text{ or } (3,1); \\ -1 & \text{otherwise.} \end{cases}$$

We define the *value* of a 3-colouring c as

$$\phi(c) = \sum_{i=1}^{n-1} w(e_i) z_c(e_i).$$

We claim that $|\phi(c_1) - \phi(c_2)| \leq 2$ for any two 3-colourings c_1 and c_2 of P that are adjacent in the graph R_P^3, i.e., that differ on one vertex of P. This is easy to check, but we give the details for completeness.

Note first that $z(e)$ changes sign if we change the colour of exactly one end vertex of e or if we exchange the colours of e. Let v_k be the (unique) vertex on which c_1 and c_2 differ, and suppose $c_1(v_k) = x$ and $c_2(v_k) = y \neq x$. If z is the unique colour that is not x or y, then the vertices v_{k-1}, v_k, v_{k+1} (when they exist) are coloured z, x, z by c_1 and z, y, z by c_2. From this we deduce that

$$z_{c_1}(e_{k-1}) = -z_{c_2}(e_{k-1}) = -z_{c_1}(e_k) = z_{c_2}(e_k), \tag{1}$$

ignoring any terms that are not defined. If $k \neq 1, n$ then

$$\begin{aligned} \phi(c_1) - \phi(c_2) &= \sum_{i=k-1}^{k} w(e_i)(z_{c_1}(e_i) - z_{c_2}(e_i)) \\ &= 2 z_{c_1}(e_{k-1}) \big(w(e_{k-1}) - w(e_k)\big), \end{aligned}$$

where the last line follows from (1). Taking the absolute value of both sides (and noting that $|w(e_{k-1}) - w(e_k)| \leq 1$) proves the claim. If $k = 1, n$, then excluding the appropriate terms from the above calculation (and noting that $w(e_1) = w(e_{n-1}) = 1$) also yields $|\phi(c_1) - \phi(c_2)| \leq 2$.

We now let c_1 be the 3-colouring that colours $v_1, v_2, v_3, v_4, \dots$ by colours $1, 2, 3, 1, \dots$, respectively, and we let c_2 be the 3-colouring that colours

$v_1, v_2, v_3, v_4, \ldots$ by colours $3, 2, 1, 3, \ldots$, respectively. Then

$$\phi(c_1) = -\phi(c_2) = \sum_{i=1}^{n-1} w(e_i) = \left\lfloor \frac{n}{2} \right\rfloor \left\lceil \frac{n}{2} \right\rceil \geq \frac{1}{4}(n^2 - 1).$$

In order to get from c_1 to c_2, the value of the colouring must necessarily change by $|\phi(c_1) - \phi(c_2)| \geq \frac{1}{2}(n^2 - 1)$. Hence, the number of recolourings required is at least $\frac{1}{4}(n^2 - 1) = \Theta(n^2)$ because each recolouring changes the value by at most 2. This completes the proof of Theorem 11. $\square$

We now generalise Theorem 11. Recall that every k-colourable chordal graph is k-colour dense by Theorem 6.

Theorem 12. *For each fixed $k \geq 2$ and each $n \geq k$, there is a k-colourable chordal (hence k-colour-dense) graph $G_k(n)$ on n vertices that has $(k+1)$-colour diameter $\Theta(n^2)$.*

Proof. The case $k = 2$ follows from Theorem 11. Assume that $k \geq 3$ and set $n' = n - k + 2 \geq 2$. Let $G_k(n)$ be the graph obtained from a path P on n' vertices $v_1, \ldots, v_{n'}$ by adding a clique on $k - 2$ new vertices $w_1, \ldots, w_{k-2}$ and inserting an edge between each v_i and each w_j. Because we can obtain $G_k(n)$ by repeatedly adding vertices adjacent to all existing vertices, $G_k(n)$ is chordal. Clearly $G_k(n)$ is k-colourable. We now show that the k-colour diameter of $G_k(n)$ is $\Theta(n'^2) = \Theta(n^2)$.

Let c_1 be a colouring of $G_k(n)$ in which the colours 1, 2 and 3 cycle on the vertices of P. Let c_2 be the colouring closest to c_1 in $R^k_{G_k(n)}$ in which only 2 colours are used on P. To recolour from c_1 to c_2 only involves recolouring vertices on P since as long as there are 3 colours used on the path, we cannot recolour any vertex not in the path. Moreover only the colours 1, 2 and 3 are available to use on the path. So we can forget about the clique and think only about the distance between the restriction to P of c_1 and c_2 in R^3_P. Using the ideas of the proof of Theorem 11, we note again that the value of c_1 is $\Theta(n'^2) = \Theta(n^2)$ and see that if P has an even number of edges the value of c_2 is 0 (else consider instead $P - v_1v_2$). As again each recolouring changes the value by at most 2, we are done. $\square$

6 Future Work

We finish our paper by posing the following two open questions.

1. Do graphs of treewidth at most $k - 1$ have a quadratic $(k+1)$-colour diameter?

2. Do k-colourable perfect graphs have a quadratic $(k+1)$-colour diameter?

We note that k-colourable chordal graphs have treewidth at most $k-1$. Also recall that chordal graphs are perfect. Hence, affirmative answers to questions 1 and 2 would form two natural generalizations of our result for the class of chordal graphs. We could pose the same two questions even after relaxing the definition of a quadratic $(k+1)$-diameter by asking for a $(k+1)$-diameter that instead of $O(n^2)$ is at most $f(k)O(n^2)$ where f is a function that only depends on k.

References

[1] D. Achlioptas, A. Coja-Oghlan, F. Ricci-Tersenghi (2011). On the solution-space geometry of random constraint satisfaction problems. *Random Structures and Algorithms* **38** (2011) 251–268.

[2] P. Bonsma, Shortest path reconfiguration is PSPACE-hard, manuscript, arXiv:1009.3217, 2010.

[3] P. Bonsma and L. Cereceda, Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances, *Theoret. Comput. Sci.* **410** (2009) 5215–5226.

[4] L. Cereceda, PhD Thesis.

[5] L. Cereceda, J. van den Heuvel, and M. Johnson, Mixing 3-colourings in bipartite graphs, *European J. Combin.* **30** (2009) 1593–1606.

[6] L. Cereceda, J. van den Heuvel, and M. Johnson, Finding Paths Between 3-Colourings, *J. Graph Theory* **67** (2011) 69–82.

[7] R. Diestel, *Graph Theory.* Springer-Verlag, Electronic Edition, 2005.

[8] G.A. Dirac, On rigid circuit graphs. Anh. Math. Sem. Univ. Hamburg **25** (1961) 71–76.

[9] M. C. Golumbic, *Algorithmic Graph Theory and Perfect Graphs*, Annals of Discrete Mathematics Vol 57, Elsevier B.V., Amsterdam, 2004.

[10] P. Gopalan, P.G. Kolaitis, E.N. Maneva, and C.H. Papadimitriou, The connectivity of boolean satisfiability: computational and structural dichotomies, *SIAM J. Comput.* **38** (2009) 2330–2355.

[11] P.L. Hammer, F. Maffray, and M. Preismann, A characterization of chordal bipartite graphs, tech. rep. (Rutgers University, New Brunswick, NJ, 1989).

[12] F. Gavril, The intersection graphs of subtrees in trees are exactly the chordal graphs, *J. Combin. Theory Ser. B* **16** (1974) 47–56.

[13] T. Ito, E.D. Demaine, N.J.A. Harvey, C.H. Papadimitriou, M. Sideri, R. Uehara, and Y. Uno, On the Complexity of Reconfiguration Problems, *Theoret. Comput. Sci.* **412** (2010) 1054–1065.

[14] M. Kaminski, P. Medvedev, M. Milanic, Shortest paths between shortest paths and independent sets, *Proceedings of the 21st International Workshop on Combinatorial Algorithms* (IWOCA 2010), Lecture Notes in Computer Science **6460** (2010) 56–67.

[15] M.J. Pelsmajer, J. Tokazy and D.B. West, New proofs for strongly chordal graphs and chordal bipartite graphs. Unpublished Manuscript, 2004.

[16] R. Uehara, Linear time algorithms on chordal bipartite and strongly chordal graphs, *Proceedings of the 29th International Colloquium on Automata, Languages and Programming* (ICALP 2002), Lecture Notes in Computer Science **2380** (2002) 993–1004.