hexagon:2611.00260v2

The Cereceda (2007) conjecture holds for bipartite 2-degenerate graphs

Dave Goel

First version: October 23, 2026. This version: November 26, 2026.

Abstract

For a graph G and a positive integer k, the k-recolouring graph $\mathcal{R}_k(G)$ has the proper k-colourings of G as vertices, two of them adjacent when they differ at exactly one vertex. A graph is d-degenerate if each of its non-empty subgraphs has a vertex of degree at most d. The Cereceda (2007) conjecture states that, for fixed d and $k \geq d+2$, the diameter of $\mathcal{R}_k(G)$ is $O(n^2)$ for every d-degenerate graph G on n vertices. For $d = 2$, the conjecture was known to hold for $k \geq 5$ by a result of Cereceda; for $k = 4$, the best bound known for all 2-degenerate graphs is $O(n^3)$, due to Bousquet and Heinrich. We prove that the diameter of $\mathcal{R}_4(G)$ is at most $2n(n+1)$ for every bipartite 2-degenerate graph G on n vertices; hence the conjecture holds for these graphs. We also show that the order n^2 cannot be improved for this class, even at girth 6. For the upper bound, we fix an ordering of the vertices in which each vertex has at most two later neighbours. In n stages, we transform any proper 4-colouring into the colouring in which every vertex has the least colour not used by its later neighbours, recolouring each vertex at most twice per stage.[1] These results are verified in Lean.[2]

1 Introduction

We prove the Cereceda (2007) conjecture for bipartite 2-degenerate graphs. The conjecture concerns *d-degenerate* graphs, in which every non-empty subgraph has a vertex of degree at most d. It asserts that, for fixed d and every number of colours $k \geq d+2$, any proper k-colouring of a d-degenerate graph on n vertices can be transformed into any other by $O(n^2)$ recolourings of single vertices, with every intermediate colouring proper. Since Cereceda settled the case $k \geq 5$ [Ce, Thm. 5.22], only the case $k = 4$ is open for 2-degenerate graphs. The conjecture is featured in van den Heuvel's survey [vdH, Conj. 3.2], and Feghali, Johnson and Paulusma [FJP] call this case "a very challenging problem". We show that any proper 4-colouring of a bipartite 2-degenerate graph on n vertices can be transformed into any other in this way by at most $2n(n+1)$ recolourings. We also show that the order n^2 of this bound cannot be improved, even at girth 6.

All graphs are finite and simple. For a graph G on n vertices and a positive integer k, the *k-recolouring graph* $\mathcal{R}_k(G)$ has the proper k-colourings of G as vertices, two of them adjacent if they differ at exactly one vertex, that is, by a *move* recolouring one vertex; we write $\operatorname{diam} \mathcal{R}_k(G)$ for its diameter. The Glauber dynamics, a Markov chain that recolours one vertex at a time, can reach every k-colouring from every other exactly when $\mathcal{R}_k(G)$ is connected [DFFV, Ce]. The *degeneracy* of a graph is the least d for which it is d-degenerate. A d-degenerate graph G has a

[1] Expanded version, with further results [G]: https://doi.org/10.5281/zenodo.23603237

[2] Thm. 1.2 and Prop. 1.3 use only Lean 4's three standard axioms (checked with `lake comparator`): https://github.com/david44hello/paper_cereceda-2007-conjecture-bipartite-2-degenerate/tree/v5.44/lean

proper $(d+1)$-colouring, and $\mathcal{R}_k(G)$ is connected for every $k \geq d+2$ (Dyer, Flaxman, Frieze and Vigoda [DFFV]; Cereceda, van den Heuvel and Johnson [CvdHJ1]; see also [Ce, Thm. 2.7]). The graph $\mathcal{R}_{d+1}(G)$ can be disconnected (for example, for $G = K_{d+1}$), and for $d \geq 3$ it can have a component of diameter $2^{\Omega(\sqrt{n})}$ [BC, Thm. 1.4], [Ce, Thm. 5.2 and p. 109]. The Cereceda (2007) conjecture is the following.

Conjecture 1.1 (Cereceda [Ce, Conj. 5.21])**.** *For all integers $d \geq 1$ and $k \geq d+2$, there is a constant C such that* $\operatorname{diam} \mathcal{R}_k(G) \leq Cn^2$ *for every d-degenerate graph G on n vertices.*

Cereceda states it with d equal to the degeneracy of G. The two forms are equivalent for graphs with at least one edge, since a d-degenerate graph with an edge has degeneracy between 1 and d, and we can take for C the largest of the constants that Cereceda's form gives for these degeneracies and the same k.

Theorem 1.2. *For every bipartite 2-degenerate graph G on n vertices,* $\operatorname{diam} \mathcal{R}_4(G) \leq 2n(n+1)$.

Proposition 1.3. *The order n^2 in Theorem 1.2 cannot be improved. For every $n \geq 7$ there is a planar bipartite graph of degeneracy 2 on n vertices whose 4-recolouring graph has diameter at least $(n-7)^2/8$, and for every $n \geq 33$ there is a bipartite graph of degeneracy 2 and girth 6 on n vertices whose 4-recolouring graph has diameter at least $(n-33)^2/128$.*

Table 1: Our results and the best previous results for $d = 2$ and four colours.

Best result before this paper	New result
$O(n^3)$ for every 2-degenerate graph [BH, Thm. 1]. Quadratic or better, to our knowledge, only for the classes (i)–(v) below, which do not cover all bipartite 2-degenerate graphs.	$\operatorname{diam} \mathcal{R}_4(G) \leq 2n(n+1)$ for every bipartite 2-degenerate graph G on n vertices (Theorem 1.2).
Order n^2 on fans, which are 2-degenerate but not bipartite [BJLPP, Thm. 12].	Order n^2 is best possible for bipartite 2-degenerate graphs, even at girth 6 (Proposition 1.3).

Cereceda proved that $\operatorname{diam} \mathcal{R}_k(G) = O(n^2)$ for every d-degenerate graph G on n vertices and every $k \geq 2d+1$ [Ce], which covers $d = 1$. With $k \geq 2d+2$ colours the diameter is linear, by a theorem of Bousquet and Perarnau [BP, Thm. 1]. The order n^2 is best possible for every d when $k = d+2$. Take a path and join every vertex of a clique on $d-1$ vertices to every vertex of the path (for $d = 2$, a *fan*). This graph is d-degenerate, and its $(d+2)$-recolouring graph has diameter of order n^2 [BJLPP, Thms. 11 and 12]. For $d+2$ colours and $d \geq 2$, the first polynomial bound, $\operatorname{diam} \mathcal{R}_{d+2}(G) = O(n^{d+1})$ for every d-degenerate graph G on n vertices, is due to Bousquet and Heinrich [BH, Thm. 1] (see Table 1 and Figure 1 for $d = 2$). The same theorem gives the bound $O(n^2)$ for $k \geq \frac{3}{2}(d+1)$. For planar bipartite graphs, which are 3-degenerate by Euler's formula, they also proved that $\operatorname{diam} \mathcal{R}_5(G) = O(n^2)$ [BH, Thm. 2]. They describe this as proving the conjecture for planar bipartite graphs. Cranston and Mahmoud noted that this result settles the case of degeneracy 3 [CM].

Since d in Conjecture 1.1 may be any upper bound on the degeneracy (and planar bipartite graphs can have degeneracy 2), the conjecture for planar bipartite graphs also asks for four colours when $d = 2$. For planar bipartite graphs, Theorems 1 and 2 of [BH] give $\operatorname{diam} \mathcal{R}_k(G) = O(n^2)$ for every $k \geq 5$, and Cereceda's theorem for $k \geq 2d+1$ [Ce] covers forests ($d = 1$, $k \geq 3$). Together with Theorem 1.2, which settles the remaining case $d = 2$, $k = 4$, this proves the conjecture for all planar bipartite graphs.

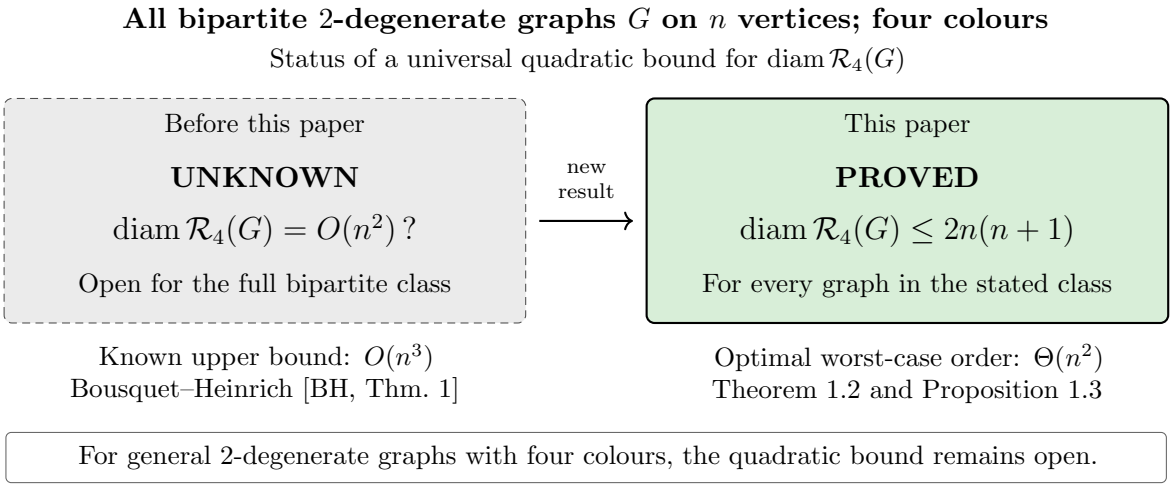

Figure 1: The universal quadratic bound for four-colour recolouring of bipartite 2-degenerate graphs. Before this paper the quadratic bound was open for the full class; the known universal upper bound was $O(n^3)$ [BH, Thm. 1]. Theorem 1.2 proves the explicit bound $2n(n+1)$. Together with the quadratic lower-bound families in Proposition 1.3, this gives optimal quadratic order in the worst case. Grey/dashed and green/solid indicate the status of the quadratic assertion; they do not denote the colours assigned to graph vertices. Here $\mathcal{R}_4(G)$ has the proper four-colourings of G as vertices, adjacent when they differ at one vertex of G.

Outline of the proofs. Section 2 reduces the sweep to a list-colouring problem on one cone; Section 3 proves the bipartite case, and Section 4 shows that the quadratic order cannot be improved. Let G be a bipartite 2-degenerate graph. Fix an ordering $v_1, \dots, v_n$ of its vertices in which every vertex has at most two later neighbours (a degeneracy ordering, see Section 2). We use the colours $0 < 1 < 2 < 3$. A vertex is *on* y if it has the colour y. It is *greedy* if it has the least colour that none of its later neighbours has. In the *sweep*, we start from any colouring and make $v_1, v_2, \dots, v_n$ greedy in turn, one per stage. Making v_t greedy may force us to move earlier vertices too, since an earlier neighbour of v_t may be on the colour that v_t must take. But only the *cone* of v_t is involved, formed by v_t and the earlier vertices from which v_t can be reached by repeatedly passing to a later neighbour. Since no vertex outside the cone has a later neighbour in it, recolouring the cone cannot make a vertex outside it non-greedy. At the start of stage t, every vertex of the cone except possibly v_t is greedy. We must reach the *target colouring* of the cone, in which every vertex is greedy.

If v_t must take a colour other than 0, the current and the target colouring put the same vertices of the cone on 0, and we leave these vertices alone. Every other vertex of the cone has a later neighbour on 0 and so at most one later neighbour not on 0. The vertices not on 0 therefore form a forest. The two colourings differ only on the tree of this forest that contains v_t. They agree on every other component. In each other component, the root (the vertex with no later neighbour in that component) has the same greedy colour in both colourings. The colours of its descendants are then forced. On the tree containing v_t, both use only the colours 1, 2 and 3. On a tree, a 3-colouring c lifts to a *height function*, an integer function h with $h \equiv c \pmod 3$ that changes by ± 1 along every edge. A move changes h by ± 2 at the moved vertex only. Two colourings whose height functions h and h' differ by an even number at every vertex are joined by moves in which each vertex v moves exactly $|h'(v) - h(v)|/2$ times. Away from v_t, both colourings alternate between 1 and 2 along the tree, so each height function h stays between $h(v_t) - 1$ and $h(v_t) + 2$. We can thus choose the two height functions to differ by 0, ± 2 or ± 4 at every vertex, and then each vertex moves at most twice.

If v_t must move to the colour 0, the vertices on 0 change. Here we use that the cone, like G, is bipartite. We construct a colouring μ of the cone that is at least as large as both the current and the target colouring at every vertex. If we raise the vertices to their colours in μ in decreasing order of μ, and then lower them to their target colours in increasing order of μ, every intermediate colouring is proper and each vertex moves at most twice. We choose μ greedily, from v_t down the ordering: each vertex takes the least colour that is at least its current and its target colour and differs from the colours of its later neighbours, read in μ inside the cone and in the current colouring outside it. The rule can fail only at a vertex on 2 in the current or target colouring. Its two later neighbours must both lie in the cone and take the colours 2 and 3 in μ.

Bipartiteness prevents this. We show that the vertices that leave 0 or go from 2 to 1 during the stage lie on one side of the bipartition, and those that join 0 or go from 1 to 2 lie on the other. Suppose that a vertex where the rule could get stuck is on 2 in the current colouring (the case of the target colouring is symmetric). Its two later neighbours are on 0 and on 1, and they cannot both change colour. Otherwise the one on 0 would leave 0 and the one on 1 would join 0 or go to 2 (a greedy vertex is never on 3). This would put them on different sides, although both are adjacent to the vertex. So one of them keeps its colour, 0 or 1. It keeps this colour in μ too, leaving 2 or 3 for the vertex. Without bipartiteness, such a μ need not exist (Remark 3.9).

In both cases we move each vertex of the cone at most twice in a stage. As the cone of v_t has at most t vertices, the sweep makes at most $2(1 + 2 + \cdots + n) = n(n + 1)$ moves. It ends at the colouring in which every vertex is greedy. This colouring is unique, since its colours are forced in the order $v_n, \ldots, v_1$. Joining two colourings through it, we obtain the bound $2n(n + 1)$.

Our sweep follows the scheme of Bonamy and Bousquet for graphs of bounded treewidth [BB], in which the vertices are settled one at a time, every earlier vertex moves at most twice per stage, and any two colourings are joined through canonical colourings, unique up to a permutation of the colours. Settling the vertices along a degeneracy ordering goes back to Cereceda's proof of the bound for $k \geq 2d + 1$ [Ce], which settles them in the opposite order and never recolours a settled vertex. In the proof of [BH, Lemma 9], Bousquet and Heinrich go in our order, $v_1, v_2, \ldots$, a fixed number of vertices at a time, and recolour only vertices already treated. From them we also take the restriction of a stage to an induced subgraph, here the cone, in which each vertex gets as its list the colours not used by its neighbours outside the subgraph. Following them, we call a colour *full* if every vertex whose list contains it has it or has a later neighbour with it; removing the vertices of a full colour lowers the degeneracy by one, as they observe. When v_t must take a colour other than 0, the colour 0 is full in both colourings. The vertices not on 0 then form a forest. For $d = 2$, Bousquet and Heinrich bound the moves for each new vertex by order n^2, which gives their $O(n^3)$. Our new ingredient is the colouring μ of Definition 3.7, with which a stage inside a bipartite cone needs at most two moves per vertex.

The lower bounds adapt two proofs of [BJLPP], one for 3-colourings of a path, which builds on [CvdHJ2], and one for 4-colourings of fans; we measure distances between colourings with height functions. We replace the apex of the fan, its vertex adjacent to all path vertices, by two additional vertices a and b, adjacent to the path vertices of odd and of even index respectively. For girth 6 we use four, $a_0, \ldots, a_3$, where a_j is adjacent to the path vertices whose index is congruent to j modulo 4. In both families, unlike the apex, these vertices keep the graph bipartite. In the two colourings that we compare they have the colour 3, and while they keep it the path cannot use it. The other three colours increase cyclically along the path in one of these colourings and decrease in the other, so if the path has ℓ vertices, its height function changes from end to end by $\ell - 1$ in one and by $-(\ell - 1)$ in the other. An additional vertex can leave the colour 3 only if its path neighbours miss one of the other three colours. In such a colouring, the absolute end-to-end change is at most 4 in the planar family and at most $3(\ell - 1)/4 + 7$ in the family of girth 6 (Lemma 4.7). A height

function with this bounded end-to-end change differs from those of the two colourings by order ℓ^2, summed over the vertices of the path, while a move changes a height function at only one vertex, by ± 2. So order ℓ^2 moves precede the first move of an additional vertex, and as many follow the last one. If the additional vertices never leave the colour 3, the moves change the height function from that of one colouring to that of the other, which also takes order ℓ^2 moves.

Comparison with known classes. Among bipartite 2-degenerate graphs, quadratic or better bounds for four colours were known, to our knowledge, only for the following five classes, with the bounds shown:

(i) graphs of treewidth at most 2, with $2n(n+1)$ [BB, Thm. 2];

(ii) graphs of maximum degree at most 3, with $O(n^2)$ [FJP, Thm. 5], and $O(n)$ for connected graphs of maximum degree 3 [BFHR, Thm. 1.3];

(iii) 3-colour-dense graphs, with $2n^2$ [BJLPP, Thm. 2];

(iv) graphs of Grundy number at most 3, with $4\chi(G)\,n$ [BB, Thm. 1], where $\chi(G)$ is the chromatic number; for bipartite G this is at most $8n$;

(v) graphs of maximum average degree (the largest average degree of a subgraph) at most $3-\varepsilon$ for a fixed $\varepsilon > 0$, with $O(n \log^2 n)$ [Feg, Thm. 1].

In (ii), Theorem 1.3 of [BFHR] assumes that, in each of the two colourings, some vertex can be recoloured. This holds for every non-empty 2-degenerate graph, since such a graph has a vertex of degree at most 2, and with four colours such a vertex can always be recoloured.

The *Grundy number* in (iv) is the largest number of colours that a *first-fit colouring* can use. A first-fit colouring takes the vertices one at a time, in some order, and gives each the least positive integer not used by an earlier neighbour. Class (iv) contains many bipartite 2-degenerate graphs, for example every graph obtained from a graph by subdividing each of its edges once. In a first-fit colouring of such a graph, a subdivision vertex that is coloured before an original neighbour v had at most one coloured neighbour and so received the colour 1 or 2. Then v receives a colour at most 3, and so does every subdivision vertex, as it has only two neighbours.

The other classes with quadratic bounds for four colours that we know of contain no bipartite graph outside (i)–(v). For example, [BB, Thm. 4] gives a quadratic bound with $k \geq \chi(G)+1$ colours for every *distance-hereditary* graph G (a graph in which any two vertices of a connected induced subgraph are at the same distance in the subgraph as in the whole graph), including four colours when $\chi(G) \leq 3$. The 3-colourable OAT graphs of Biedl, Lubiw and Merkel have a quadratic bound for four colours [BLM, Thm. 3.3]. By Remark 4.9, the bipartite graphs in both classes lie in (iii).

The union of the classes (i)–(v) does not contain all bipartite 2-degenerate graphs. For $\ell \geq 15$, the bipartite graph J_ℓ of girth 6 in Section 4.2, a path on ℓ vertices with $a_0, \ldots, a_3$, lies in none of them (Remark 4.8). For $\ell \geq 29$ these graphs give the lower bound for girth 6 in Proposition 1.3, so on them the bound of Theorem 1.2 has the best possible order n^2.

2 Degeneracy orderings and the reduction to cones

Throughout Sections 2 and 3, G is a 2-degenerate graph on n vertices. A *colouring* of G is a proper colouring with the colours $0 < 1 < 2 < 3$. Recall that a vertex is on y if it has the colour y. A *degeneracy ordering* of G is an ordering $v_1, \ldots, v_n$ of its vertices in which every vertex v has at most two later neighbours. We call them the *out-neighbours* of v and write $N_G^+(v)$ for their set. Since G is 2-degenerate, we obtain such an ordering by repeatedly deleting a vertex of degree at most 2 and listing the vertices in the order of deletion. We orient each edge from its earlier end to its later end

and call the oriented edges *arcs*. The *greedy colour* of a vertex v in a colouring c of G is the least colour that no out-neighbour of v has. We call v *greedy* in c if $c(v)$ is its greedy colour. In stage $t = 1, \dots, n$ of the *sweep*, we make v_t greedy and keep $v_1, \dots, v_{t-1}$ greedy.

Definition 2.1. For $1 \le t \le n$, the *cone* C_t is the set of vertices from which v_t can be reached along a directed path of arcs, including v_t itself.

Since arcs point to later vertices, $C_t \subseteq \{v_1, \dots, v_t\}$. Thus $\sum_t |C_t| \le n(n+1)/2$.

At stage t of the sweep, we recolour only the vertices of C_t. This is a list-recolouring problem on C_t: a vertex of C_t may not use the colours of its out-neighbours outside C_t, and v_t may use only its current colour x and its greedy colour g (we shall move it once, from x to g). A vertex of C_t has no other neighbours outside C_t, since a vertex with an arc into C_t lies in C_t. As every vertex has at most two out-neighbours and there are four colours, every vertex of C_t other than v_t keeps at least two more colours than it has out-neighbours inside C_t. This is the condition on the lists in the next definition, where we state the problem without reference to G. The instance lists C_t in decreasing original index. Thus, for an arc $u \to v_t$, u comes before v_t in the original order and v_t comes before u in the instance order. The arc itself is unchanged.

Definition 2.2. An *instance* consists of a finite set of vertices with a linear order, the *instance order*, together with out-neighbour sets and lists of colours, given as follows. Its first vertex is the *sink* s. Every vertex v has a set $N^+(v)$ of at most two *out-neighbours*, all of them before v in the instance order (the sink has none). The *edges* of the instance are the pairs $\{v, u\}$ with $u \in N^+(v)$, and the two ends of an edge are *neighbours*. The instance is *bipartite* if these edges form a bipartite graph. It is *single-sink* if s is its only vertex without out-neighbours. Every vertex v has a list $L(v) \subseteq \{0, 1, 2, 3\}$ with $|L(v)| \ge |N^+(v)| + 2$ (a 1-feasible list assignment of [BH]), and the sink has a list $\{g, x\}$ with $0 \le g < x \le 3$.

A *colouring* of an instance gives every vertex v a colour of $L(v)$, and it is *proper* if the two ends of every edge have different colours. Two colourings α and β are defined *greedily*: $\alpha(s) = x$, $\beta(s) = g$, and every other vertex v, in instance order, takes its *greedy* colour in the instance, the least colour of its list that none of its out-neighbours has. This colour exists since $|L(v)| > |N^+(v)|$. The colourings α and β depend only on the out-neighbours and the lists, not on the instance order, and both are proper, since every vertex avoids the colours of its out-neighbours. In an instance, a *move* gives one vertex another colour of its list, one that none of its neighbours has. A *sequence from α to β* consists of finitely many moves that, applied in order, transform α into β, and its *length* is the number of its moves.

For $1 \le t \le n$, the *instance on C_t* has the vertices of C_t in decreasing order of index as its instance order and v_t as its sink. The out-neighbours of a vertex $v \in C_t$ in the instance are the vertices of $N^+_G(v) \cap C_t$. Thus every vertex comes after its out-neighbours, v_t has none, and every other vertex of C_t has one in C_t, the next vertex on a directed path of arcs to v_t. So the instance on C_t is single-sink. We choose its lists in the proof of Lemma 2.5, so that α is the colouring of C_t at the start of stage t and β is the target colouring, in which every vertex of C_t is greedy. Only the case $g = 0$, treated in Section 3, needs bipartiteness and the single-sink property.

Lemma 2.3. *Let an instance have m vertices, and let its sink have the list $\{g, x\}$ with $0 < g < x$. Then there is a sequence from α to β of length at most $2m$, in which the sink moves exactly once and every other vertex at most twice.*

To prove it, we use height functions of 3-colourings of trees [GMP], [Ce, Section 3.1]. Let T be a tree and c a proper colouring of T with colours in $\{1, 2, 3\}$. A *height function* of c is a function

$h\colon V(T) \to \mathbb{Z}$ with $h(v) \equiv c(v) \pmod 3$ for all v and $|h(u) - h(v)| = 1$ for every edge $\{u, v\}$. Since T is a tree and the colours of adjacent vertices differ by ± 1 modulo 3, such a function exists (fix h at one vertex and extend it along the edges). It is unique up to adding a constant in $3\mathbb{Z}$. As in [CvdHJ2], we compare the height functions of the current and the target colouring. A move changes the current height function by ± 2 at the moved vertex only. (In a 3-colouring of a tree, a vertex can be recoloured only when all its neighbours have the same colour, and then its height changes by ± 2.) In Lemma 2.4 we adapt [Ce, Lemma 3.2] to trees with a restricted set of colours at one vertex and count the moves exactly. Cereceda attributes that lemma to Goldberg, Martin and Paterson [GMP], and our proof follows Cereceda's.

Lemma 2.4. *Let T be a tree and r a vertex of T; the allowed colours are $1, 2, 3$ at every vertex other than r, and those of a set $P \subseteq \{1, 2, 3\}$ at r. Let c, c' be proper colourings of T with allowed colours, h a height function of c, and $E\colon V(T) \to 2\mathbb{Z}$ a function such that $h + E$ is a height function of c' and $|E(r)| \le 2$ if $|P| \le 2$. Then c can be transformed into c' by exactly $\frac{1}{2}\sum_v |E(v)|$ moves through proper colourings with allowed colours, in which every vertex v moves exactly $|E(v)|/2$ times.*

The condition on $E(r)$ makes r, which may have fewer than three allowed colours, move at most once, and then directly to $c'(r)$.

Proof. We use induction on $\sum_v |E(v)|$. If $E = 0$, then $c' = c$. Suppose that $E(v) > 0$ for some v. We choose such a v with $h(v)$ least. If a neighbour u had $h(u) = h(v) - 1$, then, since $h + E$ is a height function and $E(v) \ge 2$, $(h + E)(u) \ge (h + E)(v) - 1 \ge h(v) + 1 = h(u) + 2$, so $E(u) \ge 2$ and $h(u) < h(v)$, contradicting the choice of v. Thus $h(u) = h(v) + 1$ for every neighbour u of v. In other words, v is a local minimum of h. All its neighbours have the colour congruent to $h(v) + 1$ modulo 3. We recolour v with the colour $y \in \{1, 2, 3\}$ such that $y \equiv h(v) + 2 \pmod 3$. The colouring stays proper. The colour y is allowed at v if $v \ne r$ or $|P| = 3$. If $v = r$ and $|P| \le 2$, then $E(r) = 2$. Hence $c'(r) \equiv h(r) + 2 \pmod 3$, that is, $y = c'(r)$, an allowed colour. The new colouring has the height function $h + 2\mathbf{1}_v$, where $\mathbf{1}_v$ is the indicator function of $\{v\}$. Since $E(v) \ge 2$, the function $E - 2\mathbf{1}_v$ has absolute value $|E(v)| - 2$ at v. Thus $\sum_u |E(u)|$ drops by 2. The function $E - 2\mathbf{1}_v$ still satisfies the hypotheses (if $v = r$ and $|P| \le 2$, its value at r is 0). By induction, every vertex u moves exactly $|E(u)|/2$ times in all. If $E \ne 0$ but no $E(v) > 0$, we argue symmetrically with $E(v) < 0$ and $h(v)$ largest. □

Proof of Lemma 2.3. We show that α and β put the same vertices on 0, that the other vertices induce a forest, and that α and β differ only on the tree T of this forest that contains s. Lemma 2.4 applies to T, where both colourings use only 1, 2 and 3. Since $g > 0$, the sink is not on 0 in α or in β. A vertex $v \ne s$ is on 0 in α if and only if $0 \in L(v)$ and no out-neighbour of v is on 0 in α. The same holds for β. By induction along the instance order, $\alpha^{-1}(0) = \beta^{-1}(0)$, and we call this set U. This rule also shows that the colour 0 is full in α and in β, that is, every vertex whose list contains 0 is on 0 or has an out-neighbour on 0. Let F be the set of the other vertices.

Claim 1. Every $v \in F$ has at most one out-neighbour in F. If $0 \in L(v)$, then, since v is not on 0 in α, some out-neighbour of v is on 0 and lies in U. Otherwise $|L(v)| \le 3$, and since $|L(v)| \ge |N^+(v)| + 2$, the vertex v has at most one out-neighbour. Thus F induces a forest in which each vertex has at most one *parent* (an out-neighbour in F), since on a cycle in F the vertex that comes last in the instance order would have both its neighbours on the cycle as out-neighbours in F. All neighbours of a vertex of F other than its parent and its children (the vertices of F whose parent it is) lie in U. In particular, if a vertex outside a component of the forest induced by F is adjacent to this component, then it lies in U.

Claim 2. Every vertex v of F with a parent has $\{1, 2, 3\} \subseteq L(v)$. If $0 \in L(v)$, then v has an out-neighbour on 0 and a parent, so $|L(v)| = 4$; otherwise $L(v) = \{1, 2, 3\}$.

A vertex $w \neq s$ of F without a parent has all its out-neighbours on 0, so $\alpha(w) = \beta(w) = \min(L(w) \setminus \{0\})$. A vertex of F with a parent p is not on 0, by Claim 2 its list contains 1, 2 and 3, and its other out-neighbour, if any, is on 0. So in α and in β it takes the least colour of $\{1, 2, 3\}$ other than the colour of p. By induction, we get $\alpha = \beta$ outside the component T of the forest induced by F that contains s (the vertices of U are on 0 in both colourings, and every other vertex outside T has its parent, if any, outside T). Moreover, in each colouring $c \in \{\alpha, \beta\}$, every child of a vertex on 1 is on 2, and every child of a vertex on 2 or 3 is on 1.

As a tree, T has $|T| - 1$ edges, each joining a vertex to its parent. Since each vertex has at most one parent, the sink s is the only vertex of T without one. Every other vertex of T is reached from s by passing repeatedly from a vertex to one of its children. Along such a path the colours are $1, 2, 1, 2, \ldots$ if $c(s) = 1$, they are $2, 1, 2, 1, \ldots$ if $c(s) = 2$, and they are $3, 1, 2, 1, 2, \ldots$ if $c(s) = 3$. For a height function h of $c|_T$ and every $v \in T \setminus \{s\}$, we have

$$h(v) - h(s) \in \begin{cases} \{0, 1\} & \text{if } c(s) = 1, \\ \{-1, 0\} & \text{if } c(s) = 2, \\ \{1, 2\} & \text{if } c(s) = 3. \end{cases} \tag{1}$$

We choose height functions h_α of $\alpha|_T$ and h_β of $\beta|_T$. Since $\alpha(s) \neq \beta(s)$ are colours in $\{1, 2, 3\}$, the difference $h_\beta(s) - h_\alpha(s)$ is congruent to 2 or to -2 modulo 3. Shifting h_β by a multiple of 3, we may assume that $E = h_\beta - h_\alpha$ has $E(s) \in \{-2, 2\}$. For $v \in T \setminus \{s\}$, write $E(v) - E(s) = (h_\beta(v) - h_\beta(s)) - (h_\alpha(v) - h_\alpha(s))$. By (1), both brackets lie in $[-1, 2]$. Both also have the parity of $\mathrm{dist}_T(v, s)$, since a height function changes by ± 1 along every edge. Their difference is even and lies in $[-3, 3]$, i.e. in $\{-2, 0, 2\}$. Adding $E(s) = \pm 2$ gives

$$E(v) \in \{-4, -2, 0, 2, 4\}.$$

This also holds for $v = s$, so E is even with $|E(v)| \leq 4$ for every $v \in T$.

Since $|E(s)| = 2$, Lemma 2.4 applies to T, with $r = s$ and $P = L(s) = \{g, x\}$. Every other vertex of T has a parent in T. By Claim 2, its list contains 1, 2 and 3, the only colours the moves of Lemma 2.4 use. Every vertex outside T that is adjacent to T lies in U and is on 0. So these moves are moves of the instance, U stays on 0, and no vertex outside T moves. They turn α into β in $\frac{1}{2}\sum_{v \in T} |E(v)| \leq 2|T| \leq 2m$ moves, in which s moves once and every other vertex at most twice, since $|E(s)| = 2$ and $|E(v)| \leq 4$. □

The next lemma turns a bound for the case $g = 0$ into a bound on $\operatorname{diam} \mathcal{R}_4(G)$. The reduction applies to any class of instances closed under changing the allowed lists.

Lemma 2.5. *Let $\mathcal{K}$ be a class of instances that is closed under changing the lists (to any other lists allowed in an instance), and let $\lambda \geq 2$ be such that every instance in $\mathcal{K}$ with $g = 0$ has a sequence from α to β of length at most λm, where m is its number of vertices. Let G have a degeneracy ordering such that, for every t, the instance on C_t lies in $\mathcal{K}$ for some choice of lists (and so, by the closure property, for every choice). Then*

$$\operatorname{diam} \mathcal{R}_4(G) \leq 2\lambda \sum_t |C_t| \leq \lambda n(n+1).$$

Proof. Giving $v_n, v_{n-1}, \ldots, v_1$ in turn their greedy colours, we obtain the only colouring γ in which all vertices are greedy. This is the first-fit colouring in the order $v_n, \ldots, v_1$ with the colours $0 < 1 < 2 < 3$. Bousquet and Heinrich [BH, Section 2] note that it has a full colour (here the colour

0). We show that γ can be reached from any colouring in at most $\lambda \sum_t |C_t|$ moves. We can then join any two colourings through γ by at most $2\lambda \sum_t |C_t|$ moves.

We prove that after stage t of the sweep the vertices $v_1, \dots, v_t$ are greedy. At stage t, let c be the current colouring, $x = c(v_t)$, and g the greedy colour of v_t. If $x = g$, stage t makes no move. Otherwise we take the instance on C_t with the lists $L(v_t) = \{g, x\}$ and, for every other $v \in C_t$, $L(v) = \{0,1,2,3\} \setminus c(N_G^+(v) \setminus C_t)$. Here $g < x$, since no out-neighbour of v_t has the colour x (c is proper) and g is the least colour with this property. For every other $v \in C_t$, $|L(v)| \geq 4 - |N_G^+(v) \setminus C_t| \geq |N_G^+(v) \cap C_t| + 2$, since $|N_G^+(v)| \leq 2$. So these lists are allowed in an instance.

Recall that every neighbour of a vertex $v \in C_t$ outside C_t is an out-neighbour of v. The edges of G inside C_t are exactly the edges of the instance, and the neighbours outside C_t do not move and their colours are not in the lists. Every sequence in the instance is thus also a sequence of moves in G.

Both $c|_{C_t}$ and α give v_t the colour x, the other vertices of C_t are greedy before the stage, and the greedy colour of such a vertex in G is the least colour of its list that no out-neighbour inside C_t has. By induction along the instance order, $c|_{C_t} = \alpha$. Likewise, every vertex of C_t is greedy in the colouring that agrees with β on C_t and with c elsewhere. There v_t has the colour g, every other vertex of C_t has its greedy colour in the instance, and its out-neighbours outside C_t keep their colours. A vertex v_j with $j < t$ that lies outside C_t has no out-neighbour in C_t and stays greedy.

If $g = 0$, the condition on $\mathcal{K}$ gives at most $\lambda|C_t|$ moves. If $g \neq 0$, then $0 < g < x$ and Lemma 2.3 gives at most $2|C_t| \leq \lambda|C_t|$ moves. After stage n every vertex is greedy, i.e. the colouring is γ. Summing over t gives the bound. □

3 Proof of Theorem 1.2

By Lemma 2.5, Theorem 1.2 follows from the next theorem, a bound on the length of a sequence from α to β in bipartite single-sink instances; we make the deduction at the end of this section.

Theorem 3.1. *Let a bipartite single-sink instance have m vertices. Then there is a sequence from α to β in which the sink moves exactly once and every other vertex at most twice; in particular, its length is at most $2m - 1$.*

Throughout this section, instances are single-sink. Every vertex other than the sink then has one or two out-neighbours. A vertex with two of them has the list $\{0,1,2,3\}$, since $|L(v)| \geq |N^+(v)| + 2$. If the sink has the list $\{g, x\}$ with $g \neq 0$, Lemma 2.3 gives a sequence as in Theorem 3.1. From now on, $g = 0$. We call a colour *free* at a vertex v in a colouring if no out-neighbour of v has it. Thus $\alpha(v)$ and $\beta(v)$, for $v \neq s$, are the least free colours of $L(v)$ in α and in β.

Lemma 3.2. *For every vertex $v \neq s$, both $\alpha(v)$ and $\beta(v)$ lie in $\{0,1,2\}$.*

Proof. As $|L(v)| \geq |N^+(v)| + 2$, at least $|N^+(v)| + 1$ colours of $L(v)$ lie in $\{0,1,2\}$, and the $|N^+(v)|$ out-neighbours of v cannot use all of them. Hence, in each of α and β, some colour of $L(v) \cap \{0,1,2\}$ is free. □

We call the pair $(\alpha(v), \beta(v))$ the *pattern* of v. Since $g = 0$, the pattern of the sink is $(x, 0)$. A vertex v is *fixed* if $\alpha(v) = \beta(v)$. It is then *fixed on* the colour $\alpha(v)$. A vertex with a pattern $(0, y)$, $y \neq 0$, is a *leaver* (it leaves the colour 0), and one with a pattern $(y, 0)$, $y \neq 0$, is a *joiner* (it joins the colour 0); the sink is a joiner. We shall see that, in a bipartite instance, leavers and joiners lie on opposite sides.

Lemma 3.3. *In an instance with $g = 0$, every leaver has an out-neighbour that is a joiner, and every joiner other than s has an out-neighbour that is a leaver.*

Proof. A leaver v has $0 \in L(v)$ and is not on 0 in β. Then some out-neighbour u has $\beta(u) = 0$, and $\alpha(u) \neq 0$ because $\alpha(v) = 0$. In the same way, a joiner $v \neq s$ has an out-neighbour u with $\alpha(u) = 0$, and $\beta(u) \neq 0$ because $\beta(v) = 0$. □

Lemma 3.4. *Let μ be a proper colouring of an instance with $\mu(v) \in L(v)$, $\mu(v) \geq \alpha(v)$ and $\mu(v) \geq \beta(v)$ for every v, and $\mu(s) = x$. Then the following is a sequence from α to β in which every vertex moves at most twice and the sink exactly once: first recolour every vertex v from $\alpha(v)$ to $\mu(v)$, in decreasing order of μ, and then every vertex v from $\mu(v)$ to $\beta(v)$, in increasing order of μ (ties in any order; no move when the two colours agree).*

Proof. It suffices to check, for each edge $\{v, w\}$, the moves of v and of w. Since μ is proper, $\mu(v) \neq \mu(w)$. In particular, vertices with the same value of μ are not adjacent, so their order does not matter. Suppose that $\mu(w) > \mu(v)$. In the first part w moves first, while v is on $\alpha(v) \leq \mu(v) < \mu(w)$. When v moves, w is on $\mu(w) \neq \mu(v)$. In the second part v moves first, to $\beta(v)$, while w is on $\mu(w) > \mu(v) \geq \beta(v)$. When w moves to $\beta(w)$, v is on $\beta(v) \neq \beta(w)$. All colours used lie in the lists. The sink starts on $x = \mu(s)$ and moves only in the second part. □

For example, let the sink s have the list $\{0, 1\}$ and $x = 1$, and let one more vertex u, with $N^+(u) = \{s\}$, have the list $\{0, 1, 2, 3\}$. Then $\alpha(u) = 0$ and $\beta(u) = 1$. We can take $\mu(u) = 2$. Neither s nor u can move directly to its colour in β, the current colour of the other. So u first rises to 2, then s moves to 0, and u moves down to 1 (Figure 2).

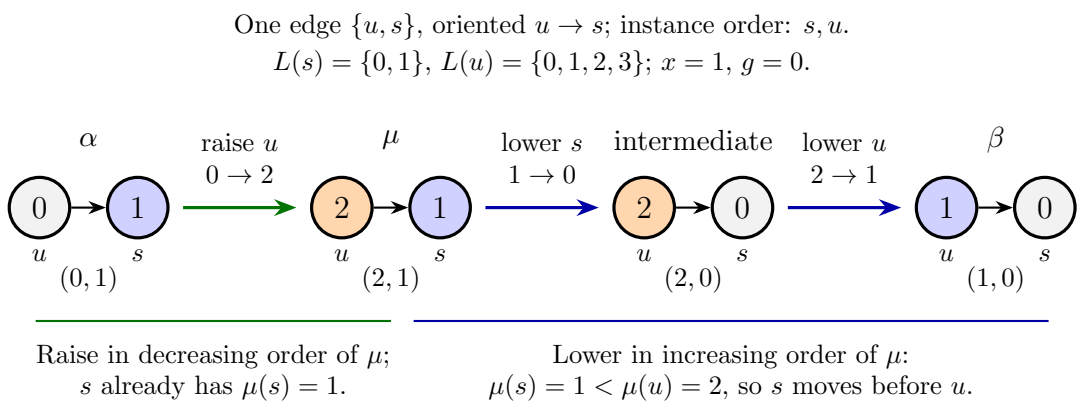

Figure 2: The upper-majorant mechanism on a bipartite single-sink instance with two vertices. The arc $u \to s$ means $N^+(u) = \{s\}$ and $N^+(s) = \varnothing$; its head comes first in the instance order. With the lists shown, the greedily defined endpoint colourings are $\alpha(u, s) = (0, 1)$ and $\beta(u, s) = (1, 0)$. Taking $\mu(s) = x = 1$ and $\mu(u) = 2$ gives a proper colouring dominating both endpoints at every vertex. Raising to μ in decreasing order of μ, then lowering to β in increasing order of μ, gives the three displayed moves; a vertex already at its required colour is skipped. Each move changes one vertex to an allowed colour different from its neighbour's colour. The numerals specify the colours independently of the shading.

It remains to find a colouring μ as in Lemma 3.4. For the rest of the section the instance is also bipartite. It is connected, since every vertex reaches s through out-neighbours. Let S_0 be the side of the bipartition that contains s, and S_1 the other side. The out-neighbours of a vertex lie on the side opposite to it.

Lemma 3.5. *In a bipartite single-sink instance with $g = 0$, each vertex other than the sink has one of the patterns listed for its side:*

$S_0 \setminus \{s\}$	S_1
$(0,0),\ (1,1),\ (2,2)$	$(0,0),\ (1,1),\ (2,2)$
$(1,0),\ (2,0),\ (1,2)$	$(0,1),\ (0,2),\ (2,1)$

The sink has pattern $(x, 0)$.

Proof. We use induction along the instance order. For a vertex v, we may assume the statement for its out-neighbours, which come earlier. By Lemma 3.2, a vertex $v \neq s$ has one of the nine patterns in $\{0,1,2\}^2$. We show that a leaver lies in S_1, a joiner other than s lies in S_0, a vertex with the pattern $(2,1)$ lies in S_1, and a vertex with the pattern $(1,2)$ lies in S_0. If v is a leaver, then by Lemma 3.3 it has a joiner out-neighbour, s or a vertex of S_0 by induction. Hence $v \in S_1$. If $v \neq s$ is a joiner, it has a leaver out-neighbour, which lies in S_1, and so $v \in S_0$. If v has the pattern $(2,1)$, then $1 \in L(v)$ since $\beta(v) = 1$. As $\alpha(v) = 2$ is the least free colour of $L(v)$ in α, the colour 1 is not free at v in α, so some out-neighbour p has $\alpha(p) = 1$, and $\beta(p) \neq 1$ since β is proper. Thus $p = s \in S_0$, or p is a joiner other than s or has the pattern $(1,2)$, and then $p \in S_0$ by induction. In either case $v \in S_1$. If v has the pattern $(1,2)$, then in the same way some out-neighbour p has $\beta(p) = 1 \neq \alpha(p)$. Here $p \neq s$ is a leaver or has the pattern $(2,1)$. By induction $p \in S_1$, and so $v \in S_0$. This excludes, on each side, exactly the three patterns not listed for it. □

We shall choose μ greedily, in instance order (Definition 3.7). A loss of all candidates can occur only when $\max(\alpha(v), \beta(v)) = 2$ and the two out-neighbours use both 2 and 3 in μ. Part (a) of the next lemma rules this out when the vertex is fixed on 2, and part (b) when it is not fixed.

Lemma 3.6. *In a bipartite single-sink instance with $g = 0$, (a) every out-neighbour of a vertex fixed on 2 is fixed on 0 or on 1; (b) if a vertex v has two out-neighbours, is not fixed, and is on 2 in α or in β, then some out-neighbour of v is fixed on 0 or on 1.*

Proof. (a) Let v be fixed on 2. If v has two out-neighbours, then $L(v) = \{0,1,2,3\}$. In α and in β, neither 0 nor 1 is free at v, so one out-neighbour of v is on 0 and the other on 1. If one of them is not fixed, then, since the two out-neighbours have the colours 0 and 1 both in α and in β, one has the pattern $(0,1)$ and the other $(1,0)$. By Lemma 3.5, a vertex with the pattern $(0,1)$ lies in S_1, and a vertex with the pattern $(1,0)$ lies in S_0 (or is $s \in S_0$). So the two out-neighbours of v would lie on different sides, which is impossible, as both lie on the side opposite to v. Hence both are fixed, on 0 and on 1. If v has one out-neighbour p, then $L(v)$, with at least three colours, contains 0 or 1. No such colour is free at v in α or in β, since 2 is the least free colour of $L(v)$ in both. So p has it in α and in β. Since p has a single colour in α, $L(v)$ contains exactly one of 0 and 1, and p is fixed on it.

(b) As v has two out-neighbours, $L(v) = \{0,1,2,3\}$. Let $c \in \{\alpha, \beta\}$ be a colouring in which v is on 2. Then neither 0 nor 1 is free at v in c, so one out-neighbour u_0 of v has $c(u_0) = 0$ and the other, u_1, has $c(u_1) = 1$. Suppose that neither is fixed. If $c = \alpha$, then u_0 is a leaver, and u_1 is $s \in S_0$ or has the pattern $(1,0)$ or $(1,2)$. If $c = \beta$, then u_0 is $s \in S_0$ or a joiner, and u_1 has the pattern $(0,1)$ or $(2,1)$. In both cases, by Lemma 3.5, u_0 and u_1 lie on different sides, which is impossible as in (a). Hence u_0 or u_1 is fixed, on 0 or on 1. □

Definition 3.7. In instance order, put $\mu(s) = x$ and, for $v \neq s$, let $\mu(v)$ be the least colour of $L(v)$ that is at least $\max(\alpha(v), \beta(v))$ and differs from $\mu(u)$ for every out-neighbour u of v.

Proposition 3.8. *In a bipartite single-sink instance with $g = 0$, μ is well defined, and $\mu(v) = \alpha(v)$ for every fixed vertex v.*

Proof. We use induction along the instance order. Let $v \neq s$ and $M = \max(\alpha(v), \beta(v))$. By Lemma 3.2, $M \leq 2$. By the definition of μ, every out-neighbour $u \neq s$ of v has $\mu(u) \geq \max(\alpha(u), \beta(u))$, and so does s if it is an out-neighbour of v, since $\mu(s) = x = \alpha(s)$. Moreover $\alpha(u) \neq \alpha(v)$ and $\beta(u) \neq \beta(v)$.

Let v be fixed on y. We show that $\mu(v) = y$, that is, that no out-neighbour u has $\mu(u) = y$. If $y = 0$, then $\mu(u) = 0$ would force $\alpha(u) = 0 = \alpha(v)$, which is impossible. If $y = 1$, then $\mu(u) = 1$ would force $\alpha(u) = \beta(u) = 0$, since both are at most $\mu(u)$ and differ from 1. Then u would be fixed on 0 (so $u \neq s$), and $\mu(u) = 0$ by induction, a contradiction. If $y = 2$, then u is fixed on 0 or 1 by Lemma 3.6(a), and then $\mu(u) \leq 1$ by induction.

Now suppose that v is not fixed. Then $M \geq 1$. We call the colours of $L(v)$ that are at least M the *candidates*. We show that some candidate differs from $\mu(u)$ for every out-neighbour u of v. If v has two out-neighbours, then $L(v) = \{0, 1, 2, 3\}$. For $M = 1$, the two out-neighbours exclude at most two of the three candidates $1, 2, 3$. For $M = 2$, v is on 2 in α or in β, and Lemma 3.6(b) gives an out-neighbour u fixed on 0 or on 1, with $\mu(u) \leq 1$ by induction. The other excludes at most one of the candidates $2, 3$. If v has one out-neighbour u, then $|L(v)| \geq 3$. For $M = 1$, only the colour 0 can lie below M, so at least two candidates remain. For $M = 2$, the colour 2 is the greedy colour of v in α or in β, that is, the least colour of $L(v)$ other than the colour of u in that colouring. So at most one colour of $L(v)$ is less than 2, since otherwise 0 and 1 would both be in $L(v)$, and one of them would differ from the colour of u. As $|L(v)| \geq 3$, $L(v)$ is $\{0, 2, 3\}$ or $\{1, 2, 3\}$, with the candidates 2 and 3. For both values of M, u excludes at most one of at least two candidates. □

Proof of Theorem 3.1. Every edge joins a vertex v to an out-neighbour u, with $\mu(v) \neq \mu(u)$ by Definition 3.7. Hence μ is proper. Also $\mu(v) \in L(v)$, $\mu \geq \alpha$, $\mu \geq \beta$ and $\mu(s) = x$. Lemma 3.4 now gives the required sequence. □

Remark 3.9. Without bipartiteness, which we used only in Lemmas 3.5 and 3.6, a colouring μ as in Lemma 3.4 need not exist. Take the sink s with the list $\{0, 1\}$ and $x = 1$, and the vertices u_1, u_2, u_3 in instance order, with $N^+(u_1) = \{s\}$, $N^+(u_2) = \{u_1\}$, $N^+(u_3) = \{u_1, u_2\}$, $L(u_1) = L(u_3) = \{0, 1, 2, 3\}$ and $L(u_2) = \{1, 2, 3\}$. This is a single-sink instance with $g = 0$, whose edges form the triangle $u_1u_2u_3$ and the edge su_1. The patterns of u_1, u_2, u_3 are $(0, 1)$, $(1, 2)$ and $(2, 0)$. Such a μ would have $\mu(u_1) \geq 1$ and $\mu(u_1) \neq \mu(s) = 1$, and $\mu(u_2), \mu(u_3) \geq 2$. So it would colour the triangle with the two colours 2 and 3, which is impossible. The conclusion of Lemma 3.6(b) fails at u_3, which has two out-neighbours, is not fixed and is on 2 in α, while neither of its out-neighbours is fixed. Still, this instance has a sequence as in Theorem 3.1, in which we recolour u_1 to 3, s to 0, u_3 to 0, u_2 to 2 and u_1 to 1. In the instance of the first open problem in Section 5, by contrast, some vertex must move three times.

Proof of Theorem 1.2. Let G be a bipartite 2-degenerate graph on n vertices. We choose a degeneracy ordering of G. The class $\mathcal{K}$ of bipartite single-sink instances is closed under changing the lists. By Theorem 3.1, it meets the condition of Lemma 2.5 with $\lambda = 2$. Every instance on a cone lies in $\mathcal{K}$: the cone induces a bipartite subgraph of G, and the instance on it is single-sink. Lemma 2.5 thus gives $\operatorname{diam} \mathcal{R}_4(G) \leq 2n(n + 1)$. □

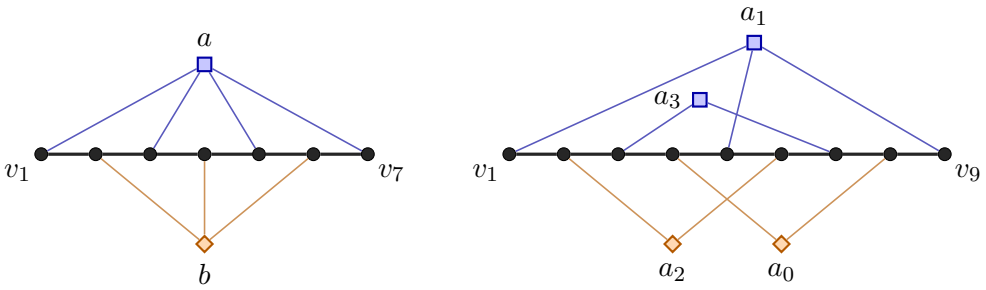

(a) H_7 (planar) (b) J_9 (girth 6)

Figure 3: Lower-bound graph families for Proposition 1.3. Dark circles and thick edges form the path spine $v_1 \cdots v_\ell$; blue squares (above) join odd-indexed path vertices, and orange diamonds (below) join even-indexed path vertices. (a) H_7, obtained by splitting the apex of a fan into a and b (Section 4.1); (b) J_9, with four auxiliary vertices $a_0, \dots, a_3$ and girth 6 (Section 4.2).

4 Proof of Proposition 1.3

We deduce Proposition 1.3 from the bounds in terms of n in Propositions 4.1 and 4.6. Here ℓ is the number of vertices of a path and n that of the whole graph. As before, a *colouring* is a proper colouring with the colours $0, 1, 2, 3$.

4.1 A planar family

For $\ell \geq 2$, let H_ℓ be the path $v_1 v_2 \cdots v_\ell$ together with a vertex a joined to every v_i with i odd and a vertex b joined to every v_i with i even. It has $n = \ell + 2$ vertices. It is planar (Figure 3 draws H_7 with a above the path and b below it), 2-degenerate (in the ordering $v_1, \dots, v_\ell, a, b$ every vertex has at most two later neighbours) and bipartite, with sides $\{v_i : i \text{ odd}\} \cup \{b\}$ and $\{v_i : i \text{ even}\} \cup \{a\}$. For $\ell \geq 3$ it contains the 4-cycle $a v_1 v_2 v_3 a$, so it is not a forest, and its degeneracy is 2.

Proposition 4.1. *For $\ell \geq 5$, the colourings σ^+ and σ^- of H_ℓ (a and b on 3, v_i on $\pm(i-1) \bmod 3$) are at distance at least $(\ell-5)^2/8$ in $\mathcal{R}_4(H_\ell)$. Hence $\operatorname{diam} \mathcal{R}_4(H_\ell) \geq (n-7)^2/8$.*

We call a colouring *plain* if a and b have the colour 3. In a plain colouring the path uses only the colours $0, 1, 2$. For a plain colouring c and $1 \leq i < \ell$, the *step* $\delta_i \in \{+1, -1\}$ is defined by $\delta_i \equiv c(v_{i+1}) - c(v_i) \pmod 3$. A *height function* of c is a function $h\colon \{1, \dots, \ell\} \to \mathbb{Z}$ with $h(i+1) - h(i) = \delta_i$ for $1 \leq i < \ell$. It is the analogue on a path of the height functions of Section 2, without the condition $h(i) \equiv c(v_i) \pmod 3$. We define the *winding* of c as $W = h(\ell) - h(1)$. It does not depend on the choice of h and is the weight of the path oriented from v_1 to v_ℓ in the sense of [CvdHJ2]. For σ^+ every step is $+1$, and for σ^- every step is -1.

Lemma 4.2. *Let two plain colourings differ by a move of a path vertex v_i, and let h be a height function of the first. Then a height function of the second is obtained from h by changing $h(i)$ by 2 or by -2, as the move determines, and keeping the other values of h.*

Proof. Suppose first that $1 < i < \ell$. The new colour of v_i is not 3 (as v_i is adjacent to a or to b, both on 3), not its old colour, and not the colours of v_{i-1}, v_{i+1}, which also differ from its old colour. As only $0, 1, 2$ are available, v_{i-1} and v_{i+1} have the same colour, so $\delta_{i-1} = -\delta_i$. After the move both steps change sign. Thus $h(i)$ changes by ± 2. For $i \in \{1, \ell\}$, the only path neighbour of v_i has the colour that differs from both the old and the new colour of v_i, so the one step at v_i changes sign, and again $h(i)$ changes by ± 2. □

Lemma 4.3. *If a plain colouring c allows a move of a or of b, then $|W| \le 4$.*

Proof. Let h be a height function of c. We call v_i an *odd* or an *even* vertex according to the parity of i. Suppose that a can move to a colour $y \in \{0,1,2\}$. Then every odd vertex has one of the other two colours, $y_1 \equiv y+1$ and $y_2 \equiv y+2 \pmod 3$. Let $u = v_{2j-1}$ and $u' = v_{2j+1}$ be consecutive odd vertices, and $w = v_{2j}$ the even vertex between them. If u and u' have the same colour, then $h(2j+1) = h(2j-1)$. Otherwise w, adjacent to u, u' and b, has neither y_1 nor y_2 nor 3 and thus has y. If u is on y_1 and u' on y_2, the colours y_1, y_1-1, y_1+1 give two steps -1, and $h(2j+1) = h(2j-1)-2$. If u is on y_2 and u' on y_1, then $h(2j+1) = h(2j-1)+2$. Along the odd vertices, changes from y_1 to y_2 and from y_2 to y_1 alternate. Hence the heights of the odd vertices take at most two values, 2 apart. The first odd vertex is v_1 and the last is v_ℓ or $v_{\ell-1}$. Thus $|h(\ell) - h(1)| \le 2+1 = 3$. If instead b can move, to a colour y, we argue in the same way with the even vertices. When two consecutive even vertices have different colours, the odd vertex between them is adjacent to both of them and to a, which is on 3; it therefore has the colour y. The first even vertex is v_2 and the last is v_ℓ or $v_{\ell-1}$. Thus $|h(\ell) - h(1)| \le 1+2+1 = 4$. □

The following count shows that a height function with small winding is far from one whose steps are all $+1$.

Lemma 4.4. *Let $h, h' \colon \{1,\dots,\ell\} \to \mathbb{Z}$ be functions whose steps $h(i+1)-h(i)$ and $h'(i+1)-h'(i)$ are all ± 1, every step of h being $+1$. Let a real number $q \ge 0$ satisfy $h'(\ell) - h'(1) \le (\ell-1) - q$ and $q/4 + 1 \le \ell/2$. Then $\sum_i |h'(i) - h(i)| \ge q^2/8$. The same holds with all signs reversed (all steps of h equal to -1 and $h'(\ell) - h'(1) \ge -(\ell-1)+q$).*

Proof. Put $f = h' - h$. The steps of f are 0 or -2. Thus f does not increase. Moreover, $f(1) - f(\ell) = (\ell-1) - (h'(\ell) - h'(1)) \ge q$. For $i \le \ell/2$, put $i' = \ell + 1 - i$. Then f decreases by at most $2(i-1)$ on $[1,i]$ and on $[i',\ell]$, so $|f(i)| + |f(i')| \ge f(i) - f(i') \ge q - 4(i-1)$. Let $\kappa = \lfloor q/4 \rfloor$. The indices $i = 1, \dots, \kappa+1$ are at most $\ell/2$, and the pairs $\{i, i'\}$ for these indices are disjoint. Summing the last inequality over them, writing $j = i-1$, and using $\kappa + 1 > q/4$ and $2\kappa \le q/2$, we obtain

$$\sum_i |f(i)| \ge \sum_{j=0}^{\kappa} (q - 4j) = (\kappa+1)(q-2\kappa) \ge (q/4)(q/2) = q^2/8. \qquad \square$$

Proof of Proposition 4.1. We consider a sequence of moves in $\mathcal{R}_4(H_\ell)$ from σ^+ to σ^-. As long as neither a nor b has moved, all colourings are plain. By Lemma 4.2, they have height functions $h_0, h_1, \dots$, where $h_0(i) = i-1$ and each h_{j+1} differs from h_j at a single index by ± 2. Each move thus changes $\sum_i |h_j(i) - h_0(i)|$ by at most 2. If a and b never move, the last height function, h', has all steps -1, so its winding differs from that of h_0 by $2\ell - 2$. We use $q = 2\ell - 4$ so that $q/4 + 1 = \ell/2$, as required by Lemma 4.4. The lemma gives $\sum_i |h'(i) - h_0(i)| \ge (\ell-2)^2/2$, so at least $(\ell-2)^2/4 > (\ell-5)^2/8$ moves are needed. Otherwise let $\tau_1 < \tau_2$ be the positions in the sequence of the first and the last move of a or b (there are at least two such moves, as σ^- is plain and each of a, b that leaves the colour 3 must return to it). Before τ_1 all colourings are plain. The colouring just before τ_1 allows a move of a or b. By Lemma 4.3 its winding is at most $4 = (\ell-1) - q$, where $q = \ell - 5$ satisfies $q/4 + 1 \le \ell/2$. Let h^* be the function h_j of this colouring. By Lemma 4.4, $\sum_i |h^*(i) - h_0(i)| \ge (\ell-5)^2/8$. Thus at least $(\ell-5)^2/16$ moves precede τ_1. Since neither a nor b moves after τ_2 and σ^- is plain, both a and b are back on 3 immediately after τ_2; every later colouring is plain. The colouring just after τ_2 allows the inverse of that last move of a or b. We read the part after τ_2 backwards from σ^-, with the height function $i \mapsto 1 - i$ of σ^- in place of h_0. Lemma 4.4 with all signs reversed then shows that at least $(\ell-5)^2/16$ moves follow τ_2. The total is at least $(\ell-5)^2/8$. □

Remark 4.5. The graph H_ℓ is the fan of [BJLPP] with its apex split into a and b. Our proof follows the proofs of [BJLPP, Thms. 11 and 12]. As in the first of them, the two colourings have the colours increasing and decreasing cyclically along the path. Where that proof bounds the change per move of the value of the colouring, a weighted sum over the edges of the path, we bound that of $\sum_i |h_j(i) - h_0(i)|$. The apex of the fan can move only when the path uses at most two colours, whereas a and b can move when the path still uses three, so we need Lemma 4.3 to bound the winding.

4.2 A family of girth 6

For $\ell \geq 5$, let J_ℓ be the path $v_1 \cdots v_\ell$ together with four vertices $a_0, \ldots, a_3$, where a_j is joined to every v_i with $i \equiv j \pmod 4$. It has $n = \ell + 4$ vertices; Figure 3 shows J_9. We obtain J_ℓ from H_ℓ by splitting a into a_1 and a_3, and b into a_0 and a_2, each new vertex keeping every second path neighbour of the old one. This removes the 4-cycles $av_iv_{i+1}v_{i+2}a$ and $bv_iv_{i+1}v_{i+2}b$. The graph J_ℓ is 2-degenerate (in the ordering $v_1, \ldots, v_\ell, a_0, \ldots, a_3$ every vertex has at most two later neighbours) and bipartite, with sides $\{v_i : i \text{ odd}\} \cup \{a_0, a_2\}$ and $\{v_i : i \text{ even}\} \cup \{a_1, a_3\}$. Its girth is 6. As the path is acyclic, every cycle contains some a_j, and the a_j are pairwise non-adjacent. A cycle through exactly one a_j is $a_jv_i \cdots v_{i'}a_j$ with $i < i'$ and $i \equiv i' \equiv j \pmod 4$. Since $i' - i \geq 4$, its length is at least 6. A cycle through two or more of the a_j has at least two path vertices between any two of the a_j that are consecutive on it (each path vertex is joined to exactly one a_j). It therefore has at least three vertices for each a_j on it, and length at least 6. Finally, the 6-cycle $a_1v_1v_2v_3v_4v_5a_1$ shows that J_ℓ is not a forest, so its degeneracy is 2.

Proposition 4.6. *For $\ell \geq 29$, the colourings σ^+ and σ^- of J_ℓ (all a_j on 3, v_i on $\pm(i-1) \bmod 3$) are at distance at least $(\ell-29)^2/128$ in $\mathcal{R}_4(J_\ell)$. Hence* $\operatorname{diam}\mathcal{R}_4(J_\ell) \geq (n-33)^2/128$.

We define plain colourings (all a_j on 3), steps, height functions and windings as for H_ℓ. Lemma 4.2 holds for J_ℓ with the same proof, since every path vertex is adjacent to some a_j.

Lemma 4.7. *If a plain colouring c of J_ℓ allows a move of some a_j, then $|W| \leq 3(\ell-1)/4 + 7$.*

Proof. Let h be a height function of c. The move goes to a colour $y \in \{0,1,2\}$ that no neighbour of a_j has; we write the other two colours as $y_1 \equiv y+1$ and $y_2 \equiv y+2 \pmod 3$. Let $p_0 < \cdots < p_r$ be the indices of the path neighbours of a_j, so that $p_{t+1} = p_t + 4$, $p_0 \leq 4$ and $p_r \geq \ell - 3$. For $0 \leq t < r$, let $D_t = h(p_{t+1}) - h(p_t)$. As a sum of four steps, D_t lies in $\{-4, -2, 0, 2, 4\}$, and $D_t \equiv c(v_{p_{t+1}}) - c(v_{p_t}) \pmod 3$. If v_{p_t} and $v_{p_{t+1}}$ have the same colour, then $D_t = 0$. A change from y_1 to y_2 gives $D_t \equiv 1 \pmod 3$ and so $D_t \in \{4, -2\}$. A change from y_2 to y_1 gives $D_t \equiv 2 \pmod 3$ and so $D_t \in \{2, -4\}$. A gap contributes $+4$ only when its endpoint colours change from y_1 to y_2. These two endpoint colours alternate whenever they change. Hence at most $\lceil r/2 \rceil$ of the D_t equal 4, and the others are at most 2:

$$-(3r+1) \leq \sum_t D_t \leq 4\lceil r/2 \rceil + 2(r - \lceil r/2 \rceil) \leq 3r + 1,$$

where the lower bound follows symmetrically. At most three steps lie before p_0, at most three after p_r, and $r \leq (\ell-1)/4$, so $|W| \leq (3r+1) + 6 \leq 3(\ell-1)/4 + 7$. □

Proof of Proposition 4.6. Put $q = (\ell-29)/4 \geq 0$. Then $3(\ell-1)/4+7 = (\ell-1)-q$ and $q/4+1 \leq \ell/2$. If no a_j moves, at least $(\ell-2)^2/4 > (\ell-29)^2/128$ moves are needed, as for H_ℓ. Otherwise we argue as in the proof of Proposition 4.1, with Lemma 4.7 in place of Lemma 4.3: Lemma 4.4 gives at least $q^2/16$ moves before the first move of a vertex a_j and, read backwards from σ^-, at least $q^2/16$ after the last, so at least $q^2/8 = (\ell-29)^2/128$ in all. □

Remark 4.8. For $\ell \geq 15$, the graph J_ℓ lies in none of the classes (i)–(v) of Section 1. (i) Its treewidth is at least 3, since it has K_4 as a minor: the sets $\{a_1, v_1, v_9, v_{10}, v_{11}\}$, $\{a_2, v_2\}$, $\{a_3, v_3\}$ and $\{a_0, v_4, \ldots, v_8\}$ are disjoint, each induces a connected subgraph, and any two of them are joined by an edge. (ii) The vertex a_1 has degree at least 4, with neighbours v_1, v_5, v_9 and v_{13}. (iv) Its Grundy number is at least 4, since a first-fit colouring in an order that begins with $v_1, v_4, v_5, v_{10}, v_7, v_8, v_9, a_1$ gives these vertices the colours $1, 1, 2, 1, 1, 2, 3, 4$. (v) As J_ℓ has $\ell + 4$ vertices and $2\ell - 1$ edges, its maximum average degree is at least $2(2\ell - 1)/(\ell + 4) > 3$.

It remains to treat (iii), whose class is defined recursively [BJLPP, Section 3]. For a vertex v, let $N(v)$ be its set of neighbours. A *separator* of a graph is a set S of vertices whose removal increases the number of components, and *identifying* two vertices u and v replaces them by one vertex adjacent to all their neighbours. A 3-colourable graph is 3-*colour-dense* if either it is a disjoint union of cliques with at most three vertices each, or it has a separator S and vertices u and v with these four properties: u and v lie in different components of the graph minus S; $S \subseteq N(v)$; the component D of u in the graph minus S satisfies $|D| = 1$ or $|D \cup S| \leq 3$; and identifying u with v gives a 3-colour-dense graph. The graph J_ℓ is connected, has more than three vertices and girth 6, and has minimum degree at least 2, since each v_i has a neighbour on the path and a neighbour a_j, and each a_j has at least two neighbours on the path. It is not a disjoint union of cliques, so it suffices to show that no S, u, v and D have the first three properties. If $|D| = 1$, then $D = \{u\}$ and $N(u) \subseteq S \subseteq N(v)$; as u has degree at least 2, u and v have at least two common neighbours, and J_ℓ has a cycle of length 4. If $|D| \geq 2$, then $|S| \leq 1$, since $|D \cup S| \leq 3$. The empty set is not a separator of a connected graph. If $S = \{z\}$, then D consists of two adjacent vertices, each of degree at least 2 with all its neighbours in $D \cup S$. So both are adjacent to z, and J_ℓ has a triangle. Both cases contradict girth 6.

Remark 4.9. The bipartite distance-hereditary graphs and the bipartite OAT graphs, named in Section 1, lie in the class (iii). Replacing 3 by 2 in the definition of 3-colour-dense graphs (Remark 4.8) gives the 2-colour-dense graphs, which are also 3-colour-dense, by induction on the number of vertices. Among the 2-colour-dense graphs are the *semi-false* graphs [BJLPP, Thm. 10], the graphs that can be built from isolated vertices by repeatedly adding a vertex adjacent to exactly one existing vertex, or a vertex u with $N(u) \subseteq N(v)$ for an existing vertex v. Semi-false graphs include the chordal bipartite graphs (bipartite graphs in which every cycle of length at least 6 has a chord) [BJLPP, Thm. 8]. A bipartite distance-hereditary graph has no induced cycle of length at least 6, since deleting one vertex of such a cycle would leave a connected induced path whose ends are at distance 2 in the graph but at least 4 in the path. Being chordal bipartite, it is semi-false and lies in (iii).

OAT graphs are built from single vertices by four operations: disjoint unions; joins (the join of two disjoint graphs adds every edge between them); adding a vertex u with $N(u) \subseteq N(v)$ for an existing vertex v; and attaching a clique Q to a vertex v (joining every vertex of Q to v). Each operation adds edges only at a new vertex or between the two graphs that it combines, so every graph built along the way is an induced subgraph of the final graph. If the final graph is bipartite, so are all of them. A join in such a construction is of two edgeless graphs, since an edge of one of them and a vertex of the other would form a triangle. So the join is a complete bipartite graph; it is chordal bipartite and therefore semi-false. Since Q together with v is a clique in a bipartite graph, an attached clique Q is a single vertex, and attaching it adds a vertex of degree 1. A disjoint union of two semi-false graphs is semi-false (start from the isolated vertices of both constructions, then perform their two sequences of operations one after the other), and adding a vertex u with $N(u) \subseteq N(v)$ keeps a graph semi-false. By induction along the construction, a bipartite OAT graph is semi-false, and so lies in (iii).

5 Discussion and open problems

Theorem 1.2 and Proposition 1.3 settle Cereceda's conjecture for all bipartite 2-degenerate graphs: the 4-recolouring diameter is at most $2n(n+1)$, and order n^2 cannot be improved even at girth 6. Together with Theorems 1 and 2 of Bousquet and Heinrich [BH] and Cereceda's theorem for $k \geq 2d+1$ [Ce], this completes the conjecture for all planar bipartite graphs. Within the cone sweep of Section 2, our proof pinpoints the exact mechanism separating bipartite 2-degenerate graphs from the general case: when the sink moves to 0, bipartiteness guarantees a proper upper majorant μ, whereas odd cycles can obstruct both the majorant (Remark 3.9) and the two-move-per-vertex bound (Problem 1 below). Four natural questions remain open.

1. Does every instance of Definition 2.2 with $g = 0$ and m vertices have a sequence from α to β of length $O(m)$? By Lemma 2.5, a positive answer would prove Conjecture 1.1 for $d = 2$ (the case $k \geq 5$ is known [Ce]). When the instance has an odd cycle, the colouring μ need not exist (Remark 3.9), and two moves per vertex are not always enough. Take the sink s with the list $\{0, 1\}$ and $x = 1$, and the vertices $u_1, \ldots, u_7$ in instance order, with the following out-neighbours and lists:

v	u_1	u_2	u_3	u_4	u_5	u_6	u_7
$N^+(v)$	$\{s\}$	$\{s\}$	$\{u_1\}$	$\{u_2, u_3\}$	$\{u_4\}$	$\{u_4\}$	$\{u_3, u_6\}$
$L(v)$	$\{0,1,2\}$	$\{0,1,2\}$	$\{0,2,3\}$	$\{0,1,2,3\}$	$\{0,1,2\}$	$\{1,2,3\}$	$\{0,1,2,3\}$

   This instance contains the 5-cycle $su_1u_3u_4u_2s$. The obstruction in [G, Remark 34.9] shows that every sequence from α to β for this instance recolours some vertex at least three times.

2. Is $\operatorname{diam} \mathcal{R}_4(G) = O(n^2)$ for every 2-degenerate graph G of girth at least 5 on n vertices? If so, the graphs of girth 6 in Proposition 1.3 show that the order n^2 is best possible. Is $\operatorname{diam} \mathcal{R}_4(G) = o(n^2)$ for 2-degenerate graphs G of girth at least 7? If, for a spacing $p \geq 2$, each a_j $(0 \leq j < p)$ is joined to the path vertices v_i with $i \equiv j \pmod p$, then the girth is at most 6 once $\ell \geq p + 2$. Indeed, if $i + p + 1 \leq \ell$ and $j \equiv i \pmod p$, then $a_j v_i v_{i+1} a_{j+1} v_{i+p+1} v_{i+p} a_j$ is a cycle of length 6, where a_p means a_0.

   For planar graphs of girth at least 7, the second question has a positive answer. Such a graph has maximum average degree less than $14/5$, and [Feg] then gives $\operatorname{diam} \mathcal{R}_4(G) = O(n \log^2 n)$. Planar graphs of girth at least 6 are 2-degenerate, and Theorem 1.2 gives $\operatorname{diam} \mathcal{R}_4(G) = O(n^2)$ for those that are bipartite. Bartier, Bousquet, Feghali, Heinrich, Moore and Pierron proved that $\operatorname{diam} \mathcal{R}_5(G) = O(n)$ for all planar graphs of girth at least 6 [BBFHMP, Thm. 1.7]. They remarked after that theorem that $\mathcal{R}_4(G)$ may also have linear diameter. The graphs J_ℓ do not rule out this linear bound, as they are not planar for $\ell \geq 9$. By Euler's formula, a planar graph of girth 6 on n vertices has at most $3(n-2)/2$ edges, while J_ℓ, of girth 6, has $\ell + 4$ vertices and $2\ell - 1 > 3(\ell + 2)/2$ edges.

3. The largest diameter $\rho(n)$ of $\mathcal{R}_4(G)$ over bipartite 2-degenerate graphs G on n vertices satisfies $n^2/8 - O(n) \leq \rho(n) \leq 2n^2 + O(n)$ (Proposition 1.3 and Theorem 1.2). What is $\lim_n \rho(n)/n^2$, if it exists?

4. For $d \geq 3$, is $\operatorname{diam} \mathcal{R}_{d+2}(G) = O(n^2)$ for every bipartite d-degenerate graph G on n vertices? For $d = 3$ this holds for planar bipartite graphs, which are 3-degenerate [BH].

Acknowledgments

This work was done overwhelmingly by AI models, steered by the author in a continuous critical dialogue, variously with Google's Gemini, Anthropic's Claude and OpenAI's ChatGPT. The author takes full responsibility for the paper.

References

Where an arXiv or open version is listed, the numbers we cite refer to it.

[BBFHMP] V. Bartier, N. Bousquet, C. Feghali, M. Heinrich, B. Moore, T. Pierron, Recoloring planar graphs of girth at least five, SIAM J. Discrete Math. 37 (2023) 332–350; arXiv:2112.00631.

[BLM] T. Biedl, A. Lubiw, O. Merkel, Building a larger class of graphs for efficient reconfiguration of vertex colouring, in: Extended Abstracts EuroComb 2021, Trends Math. 14, Birkhäuser, 2021, 286–292; full version: arXiv:2003.01818.

[BB] M. Bonamy, N. Bousquet, Recoloring graphs via tree decompositions, European J. Combin. 69 (2018) 200–213; arXiv:1403.6386.

[BJLPP] M. Bonamy, M. Johnson, I. Lignos, V. Patel, D. Paulusma, Reconfiguration graphs for vertex colourings of chordal and chordal bipartite graphs, J. Comb. Optim. 27 (2014) 132–143; open version: HAL lirmm-00782882.

[BC] P. Bonsma, L. Cereceda, Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances, Theoret. Comput. Sci. 410 (2009) 5215–5226; open version: TU Berlin DepositOnce (preprint of 18 May 2007), doi:10.14279/depositonce-14393.

[BFHR] N. Bousquet, L. Feuilloley, M. Heinrich, M. Rabie, Short and local transformations between $(\Delta+1)$-colorings, Innov. Graph Theory 2 (2025) 119–156; arXiv:2203.08885.

[BH] N. Bousquet, M. Heinrich, A polynomial version of Cereceda's conjecture, J. Combin. Theory Ser. B 155 (2022) 1–16; arXiv:1903.05619.

[BP] N. Bousquet, G. Perarnau, Fast recoloring of sparse graphs, European J. Combin. 52 (2016) 1–11; arXiv:1411.6997.

[Ce] L. Cereceda, Mixing graph colourings, PhD thesis, London School of Economics, 2007; https://researchonline.lse.ac.uk/id/eprint/131216

[CvdHJ1] L. Cereceda, J. van den Heuvel, M. Johnson, Connectedness of the graph of vertex-colourings, Discrete Math. 308 (2008) 913–919.

[CvdHJ2] L. Cereceda, J. van den Heuvel, M. Johnson, Finding paths between 3-colorings, J. Graph Theory 67 (2011) 69–82.

[CM] D. W. Cranston, R. Mahmoud, 5-Coloring reconfiguration of planar graphs with no short odd cycles, J. Graph Theory 105 (2024) 670–679; arXiv:2208.02228.

[DFFV] M. Dyer, A. D. Flaxman, A. M. Frieze, E. Vigoda, Randomly coloring sparse random graphs with fewer colors than the maximum degree, Random Structures Algorithms 29 (2006) 450–465.

[Feg] C. Feghali, Reconfiguring colorings of graphs with bounded maximum average degree, J. Combin. Theory Ser. B 147 (2021) 133–138; arXiv:1904.12698.

[FJP] C. Feghali, M. Johnson, D. Paulusma, A reconfigurations analogue of Brooks' theorem and its consequences, J. Graph Theory 83 (2016) 340–358; arXiv:1501.05800.

[G] D. Goel, Flow reconfiguration with one modulus for all graphs, and Cereceda's conjecture for bipartite 2-degenerate graphs, Zenodo (2026), https://doi.org/10.5281/zenodo.23603237

[GMP] L. A. Goldberg, R. Martin, M. Paterson, Random sampling of 3-colorings in $\mathbb{Z}^2$, Random Structures Algorithms 24 (2004) 279–302.

[vdH] J. van den Heuvel, The complexity of change, in: Surveys in Combinatorics 2013, London Math. Soc. Lecture Note Ser. 409, Cambridge Univ. Press, 2013, 127–160; arXiv:1312.2816.