Characterising graphs with no subdivision of a wheel of bounded diameter

Carmesin, Johannes

DOI:
10.1016/j.jctb.2023.01.004

Journal of Combinatorial Theory, Series B 161 (2023) 21–51

Contents lists available at ScienceDirect

Journal of Combinatorial Theory, Series B

journal homepage: www.elsevier.com/locate/jctb

Characterising graphs with no subdivision of a wheel of bounded diameter

Johannes Carmesin

University of Birmingham, United Kingdom of Great Britain and Northern Ireland

A R T I C L E I N F O

Article history:
Received 20 November 2020
Available online xxxx

Keywords:
Graph minor
Characterisation of graph classes
Shallow minors
Local separators
Graph decompositions
Duality theorem

A B S T R A C T

We prove that a graph has an r-bounded subdivision of a wheel if and only if it does not have a graph-decomposition of locality r and width at most two.

1. Introduction

In [3] we introduced local 2-separators of graphs and proved a local version of the 2-separator theorem. Here, we give an application of that theorem.

An important class of graphs is the class of series-parallel graphs. A well-known fact about this class is that a graph has a tree-decomposition of width at most two if and only if it has no K_4-subdivision. We prove an analogue of this fact in our new context of local separators.

Examples of graphs of diameter r are r-bounded subdivisions, see Fig. 1 for an example and Section 2 for details. The main result of this paper is the following.

URL: https://web.mat.bham.ac.uk/J.Carmesin/.

https://doi.org/10.1016/j.jctb.2023.01.004

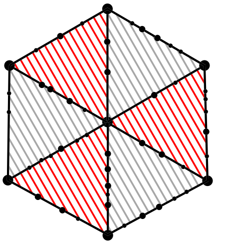

Fig. 1. An example of a subdivision of a wheel that is r-bounded. Here all the shaded triangularly shaped faces have length at most r.

Theorem 1.1. *Let G be a graph and $r \in \mathbb{N} \cup \{\infty\}$ be a parameter. Then precisely one of the following holds.*

1. *G has an r-bounded subdivision of a wheel;*
2. *G has a graph-decomposition of locality r and width at most two.*

Related results. Shallow minors were introduced in 1994 by Plotkin, Rao and Smith [13], and bounded subdivisions are a particular type of shallow minors. Recently shallow minors received attention in the context of the systematic analysis of Sparsity, see in particular the book of Nešetřil and Ossana de Mendez [12] and [7,8,10,11] for some recent work on shallow minors. While research in this area so far has focused on exploring the asymptotic relationships between various classes of sparse graphs, to my knowledge, Theorem 1.1 is the first genuine example of an explicit characterisation by excluded shallow minors or bounded subdivisions.

In [16], Thomassen proved that every graph with $e(G) \geq 2n - 2$ contains a special K_4-subdivision; that is, a cycle together with a single vertex that has three neighbours on the cycle. For $e(G) = 2n - 3$, Thomassen characterised graphs without a special K_4-subdivision in terms of admitting a $(K_3, K_{3,3})$-cockade (in modern terms these are just tree-decomposition of adhesion two with complete separators whose parts are K_3 or $K_{3,3}$).

The grid theorem [6,14] says that a graph either has a large grid as a minor or else a tree-decomposition of bounded width. In spirit this is a similar statement to Theorem 1.1: the absence of a specific substructure is characterised by the existence of a global graph-decomposition. However, there is a difference in the quantification. Indeed, Theorem 1.1 is *exact* in the sense that the two conditions in there are *mutually exclusive*, unlike for the grid theorem (even in the most recent version of [5]).

A natural continuation of our research might be to prove a grid theorem for shallow minors, as follows.

Open Problem 1.2. Given a parameter r, characterise the graphs with graph-decompositions of bounded width and locality at least r.

2. Bounded subdivisions

We assume that the reader is familiar with [3]; in particular with the definition of local cutvertices and local 2-separators; for convenience we recall the main definitions from there in Appendix A. A connected graph is *r-locally 2-connected* if it does not have an r-local cutvertex and it has a cycle of length at most r (in particular such graphs have at least three vertices). So there are no r-locally 2-connected graphs for $r < 3$. A graph is *r-locally 2-connected* if all its components are r-locally 2-connected.

A connected r-locally 2-connected graph is *r-locally 3-connected* if it does not have an r-local 2-separator and it has at least four vertices. A graph is *r-locally 3-connected* if all its components are r-locally 3-connected.

Lemma 2.1. *A 3-connected graph G whose cycles of length at most r generate all its cycles is r-locally 3-connected.*

Proof. By definition 3-connected graphs have at least four vertices.

Suppose for a contradiction that the graph G has an r-local cutvertex. Call it v. Let P be a path of G joining two neighbours of v in different components of the punctured ball $B_{r/2}(v) - v$. Let $P + v$ be the graph obtained from P by adding the vertex v together with its two incident edges to the endvertices of P. By construction $P + v$ is a cycle. By assumption the cycle $P + v$ is generated by cycles of length at most r. Pick a component of $B_{r/2}(v) - v$ containing an endvertex of the path P. Denote it by K. Let W be the set of edges incident with the vertex v to the component K. As the cycle $P + v$ intersects the edge set W oddly (in fact just once), there needs to be a generating cycle o such that it intersects the edge set W oddly; and so just once as it has no more than two edges incident with the vertex v. As o is a cycle of the ball $B_{r/2}(v)$, we conclude that $o - v$ is a subset of the component K. So o intersects the edge set W in two edges. This is a contradiction. Thus G has no r-local cutvertex.

The proof that G has no r-local 2-separator is analogous.[1] □

A *wheel* is a graph obtained from a cycle by adding a single vertex and connecting it to all vertices of the cycle. A subdivision of a wheel is *r-local* if its cycles of length at most r generate all its cycles. An *r-weighted wheel* is a graph isomorphic to a wheel with integer-valued length assigned to its edges such that if one replaced each edge by a path of its length, then the resulting graph is an r-local subdivision of a wheel.

Lemma 2.2. *r-weighted wheels are r-locally 3-connected.*

[1] The structure of that case is really the same as for local cutvertices. First suppose for a contradiction that there is a local 2-separator. Then construct a (long) cycle of G that 'traverses' this local 2-separator. By assumption it is generated by short cycles (those of length at most r). As in the other case, we conclude that there also must be a short cycle that traverses the local 2-separator. This will be the desired contradiction as local 2-separators are separating the graph locally and hence cannot be traversed by a short cycle.

Proof. Wheels are 3-connected. Hence this follows from Lemma 2.1. □

Remark 2.3. All graphs in this paper are weighted graphs; that is, graph with length assigned to their edges. All edge-lengths are positive integers. For most of our lemmas it does not make a difference that we consider weighted graphs instead of usual graphs. In the few places, where it matters, we point that out explicitly. It is essential in this paper that we work with weighted graphs as the torsos of graph-decompositions are weighted graphs.

Moreover, all graphs considered in the paper are simple; that is, they do not have loops or parallel edges.

Lemma 2.4. *Any r-local subdivision of a wheel has diameter at most r.*

Proof. We denote the center of the wheel by c. The *rim* is the unique cycle of a (subdivision of a) wheel that does not contain the vertex c. We distinguish two cases.

Case 1: the rim has length more than r. By r-locality of the wheel, every vertex is in a cycle of length at most r. This cycle must contain the center c. So every vertex has distance at most $r/2$ from the center. So the diameter is at most r.

Case 2: the rim has length at most r. Let x and y be two arbitrary vertices. By assumption they are each contained in a cycle of length at most r. If both these cycles contain the center, then both these vertices have distance at most $r/2$ from the center, so distance at most r from another. Otherwise one of these cycles must be the rim. The other cycle must contain a vertex from the rim. So x and y both have distance at most $r/2$ from that vertex, and so distance at most r from another. □

We refer to the triangles of a wheel containing the center of the wheel as the *pieces* of the wheel (these are all the triangles of the wheel except in the case of K_4). A *piece* of a subdivision of a wheel is a cycle obtained by subdividing a piece.

In our definition of when a subdivision of a wheel is r-local, we were flexible about the structure of the generating set of cycles. As most cycles in wheels use the center, we can – in fact – be more explicit as follows. We say that a subdivision of a wheel is *r-bounded* if all its pieces have length at most r. Our aim is to prove the following.

Lemma 2.5. *Every r-local subdivision of a wheel contains a subdivision of a wheel that is r-bounded.*

First we do some preparation. We say that a subdivision of a wheel has an *r-explicit* generating set if either all its pieces have length at most r or else all its pieces except for one have length at most r and the rim (that is, the unique cycle not containing the center) has length at most r.

Example 2.6. Every subdivision of a wheel with an r-explicit generating set is r-local.

Lemma 2.7. *Every subdivision of a wheel with an r-explicit generating set contains a subdivision of a wheel that is r-bounded.*

Proof. It suffices to show that every subdivision of a wheel where all but one of the pieces have length at most r and the rim has length at most r, contains an r-bounded subdivision of a wheel. In fact we will find an r-bounded subdivision of the 3-wheel K_4 as follows. By assumption there are two adjacent pieces – that is, two pieces sharing at least one edge – of length at most r. Each of these pieces shares at least one edge with the rim. Let H be the union of the rim with these two adjacent pieces. As not all edges of the rim are in these two adjacent pieces, the graph H is a subdivision of K_4. There is a unique vertex in the intersection of the old rim and the two old adjacent pieces. We see this graph $H = K_4$ as a 3-wheel with that vertex as the center. Then the pieces are precisely the two adjacent old pieces and the old rim. So the graph H is an r-bounded subdivision of a 3-wheel. □

Lemma 2.8. *Every r-local subdivision of a wheel contains a subdivision of a wheel that has an r-explicit generating set.*

In the proof, we use the following notion: a cycle is *geodesic* within a graph G if it contains a shortest path of G between any two of its vertices.

Proof. By suppressing vertices of degree two if necessary, it suffices to prove the lemma for r-weighted wheels. Let W be an r-weighted wheel. We prove this lemma by induction on the size of W (to be precise, the number of degree three vertices on its rim).

Base Case: the wheel W is the graph K_4. Assume we are given a set of cycles of length at most r that together generate all cycles of W. As every cycle of K_4 is a piece or the rim, either all cycles in the generating set are pieces or else all but one are pieces and the last one is the rim. This means that for K_4 any generating set is r-explicit.

Now assume that W is a wheel with at least four vertices on the rim.

Case 1: all geodesic cycles of W of length at most r are pieces or the rim. As every cycle in the generating set has length at most r, it is generated by geodesic cycles of length at most r. Hence the pieces and the rim together generate all cycles of W. So either all pieces of W must have length at most r, or else the rim has length at most r and all but at most one piece has length at most r. Thus the wheel W has an r-explicit generating set.[2]

Case 2: not Case 1; that is, there is a geodesic cycle o of W of length at most r that is not the rim or a piece.

Our strategy in this case will be to find an r-local subdivision of a wheel within W with less branching vertices on the rim. This smaller wheel will then have an r-explicit generating set by induction.

[2] Note that we do not apply the induction hypothesis in Case 1. So strictly speaking the base case is a special case of Case 1.

As the rim is the only cycle that does not contain the center of the wheel W, the cycle o must contain the center. As the cycle o is not a piece, it must contain at least one chord, denote it by x. Let W' be the subgraph of W obtained by deleting the chord x of the cycle o. As the wheel W is not the 3-wheel K_4, the subgraph W' is a subdivision of a wheel. Let $\mathcal{C}$ be a set of cycles of length at most r generating all cycles of W. Let P be a shortest path between the two endvertices of the chord x included in the geodesic cycle o. We obtain $\mathcal{C}'$ from $\mathcal{C}$ by replacing in each element of $\mathcal{C}$ the chord x by the subpath P of o. The set $\mathcal{C}'$ is a set of closed walks[3] of the graph W' of length at most r that generates all its cycles. Hence the subdivision W' is r-local. By induction, the graph W' has a subdivision with an r-explicit generating set. This completes the proof. □

Proof of Lemma 2.5. Combine Lemma 2.7 and Lemma 2.8. □

3. Reduction to the locally 3-connected case

A graph H is called an *r-local cut-subgraph* of a graph G if H is obtained from G by successively deleting edges and vertices and cutting at r-local 1-separators and r-local 2-separators. Here we follow the convention that after cutting at a local 2-separator, we immediately replace the torso edges by paths of the same length.

Example 3.1. A cycle of length $\ell \geq r+1$ has a path with the same number of edges as a cut-subgraph but not as a subgraph.

If H is a disconnected cut-subgraph of a graph G, then two of its edges can be copies of the same edge of G, one of which lying on a replacement path for a torso edge.

In contrast to Example 3.1, if the smaller graph H is sufficiently connected and not too big, the cut-subgraph relation is just the usual subgraph relation, as shown in the following technical lemma.

Lemma 3.2. *Assume a graph H is an r-local cut-subgraph of a graph G. If the graph H has diameter at most r, then it is a subgraph of G.*

Proof. We prove this by induction on the number of operations required to obtain H from G. Let G' be the graph obtained from G by performing the first operation. By induction, H is a subgraph of G'. If G' is a subgraph of G, we are fine. Hence we are left with two cases.

Case 1: G' is obtained from G by locally cutting a local cutvertex x. By the definition of r-local cutting, two slices of the same vertex have distance more than r. As the graph H has diameter at most r, at most one slice of the vertex x is in the graph H. If existent, denote such a slice of x in H by x'. Let G'' be the subgraph of G' obtained by deleting

[3] Every closed walk can be written as a concatenation of cycles of smaller length. Hence it suffices to construct a generating set consisting of closed walks in order to apply the induction hypothesis.

all slices of x different from x' – and if no slice of x is in H, we delete all slices of x. By construction, the graph G'' has the graph H as a subgraph. The graph G'' is equal to the subgraph of G obtained by deleting all edges incident with x that are not incident with x'. Thus H is a subgraph of the graph G.

Case 2: G' is obtained from G by r-locally cutting at an r-local separator $\{x, y\}$. By [3, Observation 5.2], no two slices of the same vertex of the local separator $\{x, y\}$ are in the graph H. If one of the vertices x or y has no slice in the graph H, then we treat this case as Case 1. Hence we may assume, and we do assume, that both vertices x and y have slices in the graph H. Denote these slices by x' and y', respectively. If these slices come from different components of $\mathrm{Ex_r}(x, y) - x - y$, then we treat each slice separately as in Case 1. Hence we may assume, and we do assume, that both slices x' and y' come from the same component of $\mathrm{Ex_r}(x, y) - x - y$. Denote that component by K.

Let P be a path of the graph G corresponding to the torso edge $x'y'$ of the graph G'. We denote the path of G' by which the torso edge $x'y'$ is replaced by P'.

If the graph H does not contain any interior vertex of the path P', then we treat each slice separately as in Case 1. Hence we may assume, and we do assume, that the graph H contains an interior vertex of the path P'.

Sublemma 3.3. *There does not exist a vertex z of G that has two copies in the graph H.*

Proof. Suppose for a contradiction there is such a vertex z. Then z must have a copy in the path P' and a copy in P. As the graph H has diameter at most r, there is a path Q of length at most r between these two copies of z in the graph H, and so in the supergraph G'. As no interior vertex of the path P' has a neighbour outside the path P', the path Q must contain one of the endvertices of the path P'. Let v be such an endvertex with $Qv \subseteq P'$. Denote by R the subpath of the path P from which the subpath Qv of P' is cloned from. Then vQR is a walk between two different slices of the vertex from which v is cloned from. This path has length at most r. As two slices never have distance at most r by [3, Observation 5.2], we derive at a contradiction. So there cannot be a vertex z of G that has two copies in the graph H. $\square$

We obtain the graph G'' from the graph G' by deleting all slices of the vertices x and y except for x' and y', deleting all torso-edge replacement paths between slices different from x' and y' – and by deleting all vertices of the paths P and P' that are not in the graph H. By construction, the graph G'' has the graph H as a subgraph.

Sublemma 3.4. *G'' is a subgraph of G.*

Proof. Label an interior vertex of the path P of G by P' if it has a clone in the graph H that is on the path P'. Label it by P if it has a clone in the graph H that is on the path P. Otherwise do not give it any label. We claim that the graph G'' is a subgraph of the graph G. To see that delete from G all edges incident with x or y that are not incident with x' or y', delete interior vertices of P without a label, for vertices with label P' delete all its incident edges that go to a vertex outside P'. By Sublemma 3.3 the resulting subgraph of G is equal to G''. □

By Sublemma 3.4, the graph G is a supergraph of the graph G'', which in turn is a supergraph of H. □

Corollary 3.5. *If a graph G contains a graph H that is an r-local subdivision of a wheel as a cut-subgraph, then H is a subgraph of G.*

Proof. By Lemma 2.4 r-local subdivisions of wheels have diameter at most r. So the corollary follows from Lemma 3.2. □

Theorem 3.6. *Let G be a graph and $r \in \mathbb{N} \cup \{\infty\}$ be a parameter. Then precisely one of the following holds.*

1. *G has an r-bounded subdivision of a wheel;*
2. *G has a graph-decomposition of locality r and adhesion at most two such that all torsos are cycles or single edges.*

Proof that Theorem 3.6 implies Theorem 1.1. The only difference between these two theorems is the formulation of the second condition. Clearly, any graph-decomposition as in condition 2 of Theorem 3.6 can be refined to yield a graph-decomposition as in Theorem 1.1. □

In the next section we prove the following.

Lemma 3.7. *Every r-locally 3-connected graph contains an r-local subdivision of a wheel.*

Proof that Lemma 3.7 implies Theorem 3.6. Let G be a graph and r be a parameter. Take the r-local block-cutvertex-graph-decomposition of G as in [3, Theorem B.1]. Then we take the r-local 2-separator decomposition [3, Theorem 1.2] of each block that is not a single edge. If all torsos of these decompositions are cycles, then we can stick these graph-decompositions together to obtain the graph-decomposition in condition 2 of Theorem 3.6.

Hence we may assume, and we do assume, that one of the torsos of these decompositions is r-locally 3-connected. Call that torso β. The 2-block γ containing β is a cut-subgraph of G. To see that cut all local cutvertices contained in that block and then delete all vertices outside that block. And β is a cut-subgraph of γ. To see that cut all local 2-separators contained in β and then delete all vertices outside β. By Lemma 3.7 the torso β has a subgraph H that is an r-local subdivision of a wheel. To summarise, H is a cut-subgraph of G. By Corollary 3.5, H is a subgraph of G; that is, G has an r-local subdivision of a wheel. By Lemma 2.5, G has an r-local subdivision of a wheel.

By Lemma 2.2 conditions 1 and 2 in Theorem 3.6 mutually exclusive. □

4. From local 3-connectivity to a bounded wheel

This section is dedicated to the proof of Lemma 3.7, which is our last step in the proof of the main result stated in the Introduction. This proof is subdivided into several subsections. In Subsection 4.1, we prove that any weighted r-locally 3-connected graph contains a geodesic cycle of length at most r with at least four edges or else a certain K_4^--subgraph. In Subsection 4.2, we use this particular K_4^--subgraph to construct a bounded wheel as a subgraph. Hence it remains to construct a bounded subdivision of a wheel using that geodesic cycle. This is done in several steps.

In Subsection 4.3 we construct a bounded fan in any r-locally 3-connected graph. In Subsection 4.4, we give conditions under which we can construct an r-local subdivision of the 3-wheel (that is, K_4) or the 4-wheel. In Subsection 4.5, and Subsection 4.6 we show how one can combine the results of the previous subsections to give a proof of Lemma 3.7.

4.1. A dichotomy result

Recall that a cycle is *geodesic* within a graph G if it contains a shortest path of G between any two of its vertices.

We say that a weighted graph G is *triangular* (with parameter r) if all its geodesic cycles of length at most r are triangles and it contains an edge e that is a shortest path between its endvertices that is in at least two triangles of length at most r.

Example 4.1. A chordal graph that has the graph K_4^- as a subgraph, where all edges have unit length, is triangular for every parameter $r \geq 1$.

Theorem 4.2. *Every r-locally 3-connected (weighted) graph has geodesic cycle of length at most r with at least four edges or it is triangular with parameter r.*

First we do some preparation.

Lemma 4.3. *Let G be a graph with a K_4^- subgraph such that its two triangles have length at most r. Then G contains a cycle of length at most r with at least four edges – or it contains an edge that is a shortest path between its endvertices that is in at least two triangles of length at most r.*

Proof. Let e be the unique edge of the K_4^--subgraph that is in both its triangles. Let P be a shortest path between the endvertices of the edge e. If the path P is equal to the edge e, we are done. Thus we may assume, and we do assume, that the path P contains at least two edges. If the path P contains at least three edges, consider the cycle $P + e$. This cycle has length at most r as it can be obtained from a triangle of K_4^- of length at most r by replacing its path of two edges between the endvertices of e by the (shortest) path P. Hence we may assume, and we do assume, that the path P consists of precisely two edges. Thus there is one of the two triangles of K_4^- of length at most r that does not contain the middle vertex of the path P. We obtain a cycle with four edges from that triangle by replacing the edge e by the path P. As P is a shortest path, this 4-cycle has length at most r. □

Lemma 4.4. *Every r-locally 3-connected graph G contains a cycle of length at most r with at least four edges – or it is triangular.*

Proof. As G is r-locally 2-connected, it includes a cycle o of length at most r. As we are done otherwise, we may assume, and we do assume, that o is a triangle. As G is r-locally 3-connected, the component containing o has another vertex. Let v be a vertex of o that has a neighbour outside o. Denote this neighbour by w. Consider the punctured ball $B_{r/2}(v) - v$. In there, there is a path P from w to some neighbour of v on the triangle o. Then $P + v$ is a cycle (recall the definition of $P + v$ from the proof of Lemma 2.1). By [3, Lemma 4.3] and [3, Lemma 3.1], the cycle $P + v$ is generated by cycles of length at most r. Similarly as above, we may assume, and we do assume, that all of them are triangles. As the cycle $P + v$ contains precisely one edge incident with the vertex v on o, one of the generating triangles must contain an odd number of edges incident with v on the cycle o. Thus there is a generating triangle o' sharing precisely one edge with the triangle o. Hence the two triangles o and o' form a K_4^- subgraph.

If G has a cycle of length at most r that is not a triangle, we are done. Otherwise, by Lemma 4.3 the graph G is triangular. □

Proof of Theorem 4.2. By Lemma 4.4, we may assume, and we do assume, that the graph G has a cycle o of length at most r with at least four edges.

Pick a cycle o of minimal length amongst all cycles with at least four edges. If o is geodesic, we are done. So suppose that the cycle o is not geodesic. Then there is a shortcut between two of its vertices; that is a path joining two vertices of the cycle o whose length is strictly shorter than the distance between these vertices on the cycle o. Such a shortcut cuts the cycle o into two cycles, each of strictly smaller length. By minimality of o, both these new cycles must be triangles. In particular, the shortcut must be a single edge. Denote that edge by e. To summarise, we have found an edge e that is a shortest path between its endvertices that is contained in two triangles of length at most r.

If all geodesic cycles of length at most r of G are triangles, then G is triangular. Otherwise G has a geodesic cycle of length at most r that is not a triangle, which is the other outcome of the theorem. This completes the proof. □

4.2. The triangular case

In this subsection we prove the following special case of Lemma 3.7.

Lemma 4.5. *Let G be a (weighted) r-locally 3-connected graph that is triangular with parameter r. Then G contains an r-weighted wheel as a subgraph.*

Proof. As G is triangular, it contains an edge e that is a shortest path between its endvertices, and there are two triangles of length at most r containing e. Denote the two endvertices of the edge e by v and w. Denote the two vertices on the triangles not incident with the edge e by x and y. To summarise, the vertices x, y, v and w span a K_4^--subgraph.

Now we construct the following auxiliary graph. Its vertex set consists of the edges incident with the vertices v or w different from the edge e. Two such edges e_1 and e_2 are adjacent in this auxiliary graph if there is a geodesic cycle of G of length at most r containing the edges e_1 and e_2. We denote this auxiliary graph by H.

Sublemma 4.6. *In the graph H there is a path from the vertex vx of H to the vertex vy.*

Proof. By assumption, the punctured explorer-neighbourhood $\mathrm{Ex_r}(v, w) - v - w$ is connected. Let P be a path of $\mathrm{Ex_r}(v, w)$ from the vertex x to the vertex y. Let o be the cycle obtained from the path P by adding the vertex v. By [3, Lemma 4.3] and [3, Lemma 3.1], there is a set $\mathcal{C}$ of cycles of length at most r of G generating the cycle o. As the geodesic cycles generate all cycles, we pick the set $\mathcal{C}$ so that it only contains geodesic cycles. As the graph G is triangular, all cycles in the set $\mathcal{C}$ are triangles.

Now let k be the connected component of the graph H containing the vertex vx. Each triangle in the set $\mathcal{C}$ containing the vertex v or w gives rise to precisely one edge of the graph H. So each of the generating triangles contains an even number of edges that are vertices of the component k. As the cycle o is equal to the sum over the cycles in $\mathcal{C}$ – evaluated over the field $\mathbb{F}_2$ –, it also must contain an even number of edges that are vertices of the component k. By construction, it contains precisely two edges that are vertices of the auxiliary graph H; these are the edges vx and vy of the graph G. And the vertex vx is in the component k by construction. So also the vertex vy of H must be in the component k. Thus the vertices vx and vy are in the same component of the graph H. □

An edge of the auxiliary graph H between vertices e_1 and e_2 is *green* if all geodesic cycles of the graph G of length at most r containing e_1 and e_2 contain the edge e. Other edges of H are *red*. We distinguish two cases.

Case 1: there is a path in the graph H from the set $\{vx, wx\}$ to the set $\{vy, wy\}$ consisting only of red edges.[4] Let $o_1, ..., o_n$ be cycles giving rise to the edges of that path. As all these edges are red, we may pick, and we do pick, the cycles o_i such that they do not contain the edge e. As G is triangular, all these cycles are triangles. So none of these triangles can use both of the vertices v and w. So there is one of these vertices, say v,

[4] This case includes the trivial case where xy is an edge of G. We will not use this in our argument.

that is a vertex of all these triangles. Thus the vertex v is the center of a wheel, whose pieces are the triangles o_1,..., o_n, vxw and vyw. As all its pieces have length at most r, this wheel is an r-weighted wheel.

Case 2: not Case 1; that is, there is a cut of the graph H separating the set $\{vx, wx\}$ from $\{vy, wy\}$ that consists only of green edges.

Let S be the set of vertices reachable from the set $\{vy, wy\}$ by a path using only red edges. Let b be the cut of the graph H between the set S and its complementary vertex set. The cut b must be nonempty by Sublemma 4.6. Let g be an arbitrary edge of the cut b. By construction of b, there is a red path (that is a path all whose edges are red) from $\{vx, wx\}$ to an endvertex of g. Denote this path by P and the endvertex of P by z. For later reference, we point out that only one endvertex of the edge g can be in the set $\{vx, wx\}$, as both these vertices are on the same side of b.

Let $o_1, ..., o_n$ be cycles giving rise to the edges of the path P. As all these edges are red, we may pick, and we do pick, the cycles o_i such that they do not contain the edge e. As G is triangular, all these cycles are triangles. So none of these triangles can use both of the vertices v and w. So there is one of these vertices, say v, that is a vertex of all these triangles. Let o be a cycle giving rise to the edge g.

We claim that the vertex v is the center of a wheel, whose pieces are the triangles o_1,..., o_n, o and vxw. As all these cycles have length at most r, this wheel would be an r-weighted wheel. So all that remains to show is that there are at least three pieces; that is, that there is at least one cycle o_i. In other words, the path P must not consists of a single vertex. This is not the case, indeed, as the triangle o would then share two edges with the triangle vxw. Then the triangles o and vxw would be identical, so the cycle o would give rise to the edge between the vertices vx and wx. As pointed out above, this violates the choice of the edge g. Hence there is at least one cycle o_i, and so o_1,..., o_n, o and vxw form the pieces of an r-weighted wheel centered at the vertex v. □

4.3. Constructing fans

A *fan* of parameter r centered around a vertex v is a sequence of oriented cycles $o_1, ..., o_n$ that all have length at most r and contain the vertex v that satisfy the following.

1. The intersection $o_i \cap o_j = \{v\}$ if $|i - j| \geq 2$; and
2. each cycle o_i has the form $vL_iM_iR_iv$ for subpaths L_i, M_i and R_i such that $R_i = L_{i+1} = o_{i+1} \cap o_i$ for all[5] $i \in [n-1]$.

See Fig. 2. We refer to the vertex v as the *center* of the fan. The *start* (or starting edge) of a fan is the first edge of the directed path L_1 and the *end* (or ending edge) of the fan is the last edge of directed path R_n. Note that the starting edge and ending edge are always incident with the center of the fan.

[5] We denote by $[n]$ the set of the first n natural numbers except zero; that is $\{1, 2, .., n\}$.

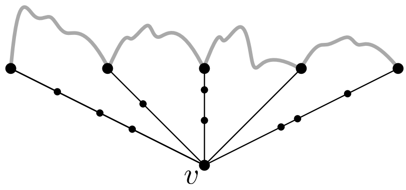

Fig. 2. A fan centered at the vertex v. The paths M_i are highlighted in grey.

A *pre-fan* of parameter r centered around the vertex v is a sequence of oriented cycles $o_1, ..., o_n$ that all have length at most r and contain the vertex v such that the vertex before after v on o_i is equal to the vertex just after v on o_{i+1} (for $i \in [n-1]$). The *start* of a pre-fan is the edge of o_1 just after v and the *end* of the fan is the edge just before v on o_n.

Example 4.7. Every fan is a pre-fan.

Roughly speaking, the next lemma gives a way how pre-fans can be 'improved to' fans.

We say that a pre-fan $(v, o_1, ..., o_n)$ *contains* another pre-fan $(v, o'_1, ..., o'_n)$ if $\bigcup o_i \supseteq \bigcup o'_i$; that is, the first pre-fan contains the second as a subgraph. Given a pre-fan $F = (o_1, ..., o_n)$, we refer to the cycles o_i as the *pieces* of the pre-fan F.

Lemma 4.8. *Every pre-fan $\mathcal{F}'$ of parameter r centered at v contains a fan $\mathcal{F}$ of parameter r centered at v that has the same start and end.*

Moreover, if the pre-fan $\mathcal{F}'$ has at least two pieces and the start is only an edge of its first cycle, and its end is only an edge of its last cycle, then the fan $\mathcal{F}$ has at least two pieces.

Proof. Let $(v, o_1, ..., o_n)$ be a pre-fan of parameter r centered at v. We consider the following operation. Assume two cycles o_i and o_j with $i < j$ share a vertex x. Let R_i be the path from x to v in the cyclic orientation of oriented cycle o_i, and let L_j be the path from v to x in the cyclic orientation of the oriented cycle o_j. We obtain o'_i from o_i by replacing the directed path R_i by the reverse of the directed path L_j. Similarly, we obtain o'_j from o_j by replacing the directed path L_j by the reverse of the directed path R_i. If the path R_i is not longer than the path L_j, then $(v, o_1, ..., o_i, o'_j, o_{j+1}, ..., o_n)$ is a pre-fan of parameter r centered at v. Otherwise $(v, o_1, ..., o_{i-1}, o'_i, o_j, ..., o_n)$ is a pre-fan of parameter r centered at v. We refer to this new pre-fan as the *reduction* of the pre-fan $(v, o_1, ..., o_n)$ at the vertex x along the indices i and j. Clearly a reduction of a pre-fan $\mathcal{F}'$ has the same start and end as $\mathcal{F}'$ and is contained in $\mathcal{F}'$.

Fans are fixed points for the reduction operation. In fact, it will follow from this proof that they are the only fixed points. This might suggest the following proof strategy.

Given a pre-fan, we shall iteratively apply reductions to it. We shall show that this procedure eventually stops, and that when it stops we have reduced the pre-fan to a fan.

Sublemma 4.9. *Assume $(v, \hat{o}_1, ...\hat{o}_m)$ is obtained from $(v, o_1, ...o_n)$ by a reduction at a vertex x along indices $i < j - 1$. Then $n > m$.*

Proof. By construction $m = n - (j - i + 1)$. □

By Sublemma 4.9, each time we perform a reduction along indices with distance at least two, the number of cycles shrinks. Hence eventually such reductions must be no longer possible; that is, cycles o_i and o_j with $i < j - 1$ can only intersect in the vertex v. For $i \in [n-1]$, let R_i be a maximal subpath of the oriented cycle o_i ending at the vertex v such that (the reverse of the directed path) R_i is a subpath of the oriented cycle o_{i+1}. Let $L_{i+1} = R_i$. Let M_i be the subpath of the oriented cycle o_i from the last vertex of the path L_i to the first vertex of the path R_i.

We have shown that the cycle o_i cannot intersect the cycle o_{i+1} in the subpath $L_i - v$ (as this would intersect the cycle o_{i-1}). If they intersect in an interior point of the path M_i, then applying a reduction at this point increases the length of the path R_i but leaves all paths R_k with $k \neq i$ invariant. Also cycle length cannot increase during such a reduction. Hence we can only perform a bounded number of such reductions. After such reductions are no longer possible, the intersection of the cycles o_i and o_{i+1} is precisely R_i. Doing this analysis for all indices $i \in [n-1]$ yields that we have a fan.

To see the 'Moreover'-part note that under these assumptions our constructions ensure that the first piece never contains the ending edge and the last piece never contains the starting edge. Thus the final fan needs to have at least two pieces. □

Roughly speaking, the next lemma gives conditions under which pre-fans exist.

Lemma 4.10. *Let G be a (weighted) r-locally 3-connected graph. Assume there is a cycle o' containing vertices v_0 and v_1 that are not adjacent on the cycle o'. Assume o' includes a shortest path P from v_0 to v_1.*

There is a pre-fan of parameter r centered at some vertex v_i none of whose pieces contains the vertex v_{i+1} (for some $i \in \mathbb{F}_2$). And there is a cycle o of length at most r including P such that the start and end of the pre-fan are on the cycle o.

Moreover, if $o \neq o'$, then o is geodesic.

Remark 4.11. If all edges of the graph G have length one, then the vertices v_0 and v_1 cannot be adjacent in the graph G. Indeed, then the path P would have length one and thus v_0 and v_1 would be adjacent on the cycle o'.

Proof of Lemma 4.10. A cycle R of the explorer-neighbourhood $\mathrm{Ex_r}(v_0, v_1)$ is *valid* if it contains the vertex v_0 but not the vertex v_1, and it contains the two neighbours of v_0 on o'.

Sublemma 4.12. *There is a valid cycle.*

Proof. As the graph G is r-locally 3-connected, the punctured explorer-neighbourhood $\mathrm{Ex}_\mathrm{r}(v_0, v_1) - v_0 - v_1$ is connected. So there is a path Q of $\mathrm{Ex}_\mathrm{r}(v_0, v_1) - v_0 - v_1$ joining the two neighbours of the vertex v_0 on the cycle o', which by assumption are both different from the vertex v_1. Then $Q + v_0$ is a valid cycle. □

A cycle is *friendly* if it has length at most r, and if it contains both vertices v_i, then it has the path P as a subpaths; moreover we require that it is geodesic.

Sublemma 4.13. *Any valid cycle has a generating set of friendly cycles.*

Proof. Let Q' be a valid cycle. By [3, Lemma 4.3] and [3, Lemma 3.1] Q' is generated by cycles of length at most r. We shall show by induction that any cycle x of length at most r has a generating set of friendly cycles. If x is not geodesic, then it is generated by two shorter cycles, which by induction have a generating set of friendly cycles. Hence we may assume, and we do assume, that the cycle x contains the vertices v_0 and v_1 but not the path P. Denote by X_1 and X_2 the two subpaths of the cycle x from v_0 to v_1. As P is a shortest path from v_0 to v_1, the closed walks PX_1 and PX_2 both have length at most the length of x. If one closed walk PX_i is not a cycle, then it is an edge-disjoint union of cycles, and all cycles in this union have strictly fewer edges than the original cycle x. If PX_i is a non-geodesic cycle, then it is generated by two shorter cycles. To summarise, either PX_i is a friendly cycle or by induction it is generated by a set of friendly cycles. Thus we have shown that any cycle x in the generating set for the valid cycle Q' has a generating set of friendly cycles, so Q' has a generating set of friendly cycles. □

By Sublemma 4.12 and Sublemma 4.13, there is a valid cycle with a generating set of friendly cycles. Call such a valid cycle R. Denote a generating set of R consisting of friendly cycles by $\mathcal{C}$.

By $\mathcal{D}$ we denote the set of all $c \in \mathcal{C}$ that contain both vertices v_0 and v_1. We distinguish two cases.

Case 0: $|\mathcal{D}|$ is even.

Now define the following auxiliary graph. Its vertex set are the cycles of $\mathcal{C} \setminus \mathcal{D}$ that contain the vertex v_0, the two subpaths of o' between v_0 and v_1, and for each $d \in \mathcal{D}$ we take the v_0-v_1-path $d \setminus P$. We denote the set of these paths by $\mathcal{D}'$. We add an edge between any two vertices of the auxiliary graph whose corresponding cycles or paths share an edge incident with the vertex v_0. We refer to this auxiliary graph as A_0.

Sublemma 4.14. *In the graph A_0, there is a path from the vertex P to some other vertex in $\mathcal{D}' \cup \{o' \setminus P\}$.*

Proof. Let K be the component of the graph A_0 containing the vertex P. Denote by e the edge of P incident with its endvertex v_0. Consider the sum S over all $c \in \mathcal{C} \setminus \mathcal{D}$ that are in the component K (over $\mathbb{F}_2$). This sum has even degree at every vertex – in particular the vertex v_0. As $|\mathcal{D}|$ is even, and the sum over all $c \in \mathcal{C}$ is nontrivial at the edge e, also the sum S is nontrivial at the edge e. So there must be another edge incident with the vertex v_0 at which the sum S is non-zero. Denote such an edge by f.

If the edge f is in the cycle R, then f is on the cycle o'. As the edge f is not on the path P, it must be on the path $o' \setminus P$. Thus the vertex $o' \setminus P$ of the graph A_0 would be in the component K, as desired. Otherwise the edge f is not on R. So it must be on some $d \in \mathcal{D}$. As the edge f is not on the path P, it must be on the path $d \setminus P$. Thus the vertex $d \setminus P$ of A_0 would be in the component K, which completes the proof. $\square$

By Sublemma 4.14, there is a path Q in the graph A_0 from P to some vertex in $\mathcal{D}' \cup \{o' \setminus P\}$. We pick a shortest such path. Then all interior vertices have their associated cycles in the set $\mathcal{C} \setminus \mathcal{D}$. These cycles form the pieces of a pre-fan centered at the vertex v_0 of parameter r. Denote that pre-fan by F. If the endvertex of the path Q is $o' \setminus P$, we pick $o = o'$. Otherwise there is some $d \in \mathcal{D}$ such that $d \setminus P$ is the endvertex of the path Q. We pick $o = d$. By construction, the path P is included in the cycle o, and this cycle has length at most r. Also the starting edge and the ending edge of the pre-fan F are on o. As none of the cycles of the pre-fan contains the vertex v_1 by construction, this completes this case.

Moreover, if $o \neq o'$, then $o \in \mathcal{C}$ and so it is geodesic.

Case 1: $|\mathcal{D}|$ is odd. This case is somewhat similar to Case 0 with the roles of 'v_0' and 'v_1' interchanged.

The graph A_1 is defined like the graph A_0 with the vertex 'v_1' in place of the vertex 'v_0'. We define $\mathcal{D}'$ as in Case 0. Similarly as Sublemma 4.14 one proves the following.

Sublemma 4.15. *In the graph A_1, there is a path from the vertex P to some other vertex in $\mathcal{D}' \cup \{o' \setminus P\}$.*

Proof. This is the same as the proof of Sublemma 4.14 with 'v_1' in place of 'v_0' and 'A_1' in place of 'A_0', and 'odd' in place of 'even'. In particular, we argue as follows: 'As $|\mathcal{D}|$ is odd, and the sum over all $c \in \mathcal{C}$ is trivial at the edge e, the sum S is nontrivial at the edge e.' Indeed, unlike in Case 0, here $e \notin R$. $\square$

The rest of this case is the same as in Case 0 with 'v_1' in place of 'v_0' and 'A_1' in place of 'A_0'. $\square$

4.4. Detecting bounded subdivisions of the 3-wheel or the 4-wheel

The purpose of this subsection is to prove Lemma 4.16 stated below.

A *theta-graph* (of parameter r) is a graph θ consisting of two vertices v and w and three internally disjoint paths between them such that at least two of the cycles obtained by concatenating a pair of these paths have length at most r. The vertices v and w are referred to as the *branching vertices* of θ. The three paths from v to w are referred to as *arms*, see Fig. 3.

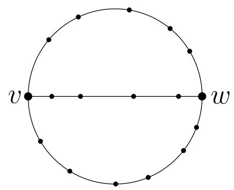

Fig. 3. A theta-graph with branching vertices v and w.

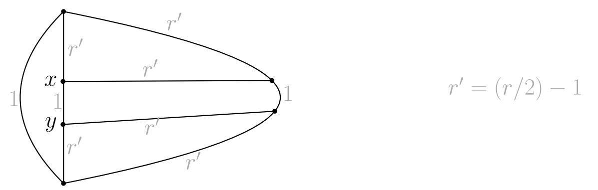

Fig. 4. A graph that is 3-connected but not r-locally 3-connected as can be seen by considering the explorer-neighbourhood $\mathrm{Ex_r}(x, y)$. Lengths of edges are given in grey.

Lemma 4.16. *Let G be a graph that contains a theta-graph θ of parameter r such that there is a cycle o of length at most r containing a branching vertex of θ and a path P between interior vertices of different arms of θ that avoids the branching vertices of θ. Then G contains an r-local subdivision of the 3-wheel or the 4-wheel.*

Example 4.17. The assumption that the cycle o contains a branching vertex cannot be omitted. An example demonstrating this is depicted in Fig. 4. Indeed, this graph has a theta-subgraph of parameter r, whose two branching vertices are denoted by x and y in the figure. This subgraph is obtained from the graph depicted by deleting the two edges of length $r' = (r/2)-1$ that are incident with neither x nor y. This graph does not contain an r-local subdivision of a wheel as the graph obtained by r-locally cutting at the r-local 2-separator $\{x, y\}$ is a series-parallel graph. Hence this graph has a graph-decomposition of locality r and width two.

Instead of proving Lemma 4.16 directly, we shall deduce it from a variant, Lemma 4.21 stated below. Next we develop the context of Lemma 4.21.

A graph G is in the class $\mathcal{W}$ if it is obtained from a theta-graph θ of parameter r by attaching[6] a path P at interior vertices of different arms of the theta-graph θ such that G contains a cycle of length at most r including P. A *weighted suppression* of a graph is obtained by iteratively suppression vertices of degree two. Here the length of the suppression edge is the sum of the length of the two edges incident with the suppressed vertex.

[6] Here *attaching* a path means that we add this path disjointly and then identify its endvertices as prescribed.

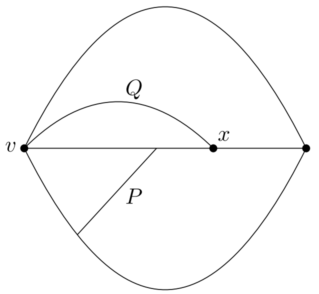

Fig. 5. A graph in the class $\mathcal{W}^*$.

Lemma 4.18. *Any graph in $\mathcal{W}$ is a subdivision of the graph K_4 such that there is a spanning tree all whose fundamental*[7] *cycles have length at most r.*

In particular, weighted suppressions of graphs in $\mathcal{W}$ are r-locally 3-connected.

Proof. Clearly any graph in the class $\mathcal{W}$ is a subdivision of the graph K_4. It remains to construct a spanning tree such that all its fundamental cycles have length at most r.

Let o_1 and o_2 be two cycles of length at most r included in the theta-graph. Let X be a subpath of the theta-graph that together with the path P forms a cycle of length at most r. Note that the path X avoids at least one edge of every arm of the theta graph. Pick an edge e_1 on an arm of the theta-graph that is in o_1 but not in o_2 so that e_1 is not in X. Similarly, pick an edge e_2 on an arm of the theta-graph that is in o_2 but not in o_1 so that e_2 is not in X. Now pick an edge e_3 on the path P arbitrarily. By construction the graph $G - e_1 - e_2 - e_3$ is a spanning tree of G. Its fundamental cycles are o_1, o_2 and XP, all of which have length at most r as $G \in \mathcal{W}$.

The 'In particular'-part follows from the fact that the fundamental cycles of any spanning tree generate all cycles and Lemma 2.1. □

A graph G is in the class $\mathcal{W}^*$ if it is obtained from a theta-graph θ of parameter r by attaching a path P at interior vertices of different arms of the theta-graph θ and another path Q at a branching vertex and some interior vertex of an arm that also contains a vertex of P such that in this arm this endvertex of P is in between the two endvertices of Q – in such a way that G contains a cycle of length at most r including both P and Q, see Fig. 5.

Lemma 4.19. *Any graph in $\mathcal{W}^*$ is a subdivision of the 4-wheel and all its cycles are generated by cycles of length at most r.*

In particular, the weighted suppressions of graphs in $\mathcal{W}^*$ are r-locally 3-connected.

[7] Given a spanning tree T and an edge e not in T, the *fundamental cycle* of e (with respect to T) is the unique cycle of the graph $T + e$.

Proof. Let G be a graph in the class $\mathcal{W}^*$. Remove the branching vertex of the theta-graph that has degree four in G. Call that vertex v. Then remove all vertices of degree one iteratively. The resulting graph is a cycle o. In the graph G, there are four paths from o to v that only intersect at the common vertex v. Thus G is a subdivision of a 4-wheel.

Let o_1 and o_2 be two cycles of the theta-graph of length at most r. Note that o_1 and o_2 together cover all edges of the theta-graph. By assumption there is a cycle u of length at most r including the paths P and Q. Let x be the endvertex of the path Q different from the branching vertex v. One of the cycles o_1 or o_2, say o_1, contains the vertex x. So the cycle o_1 contains a path of length at most $r/2$ between the vertices v and x. Denote this path by R_1. The cycle u contains a path of length at most $r/2$ between v and x. Denote that path by R_2. As the path R_2 contains precisely one of the paths P and Q, it is distinct from the path R_1. Thus the closed walk obtained by concatenating the paths R_1 and R_2 includes a cycle. Denote that cycle by o_3. See Fig. 5.

Sublemma 4.20. *The cycles o_1, o_2, o_3 and u generate all cycles of the graph G.*

Proof. The cycles o_1 and o_2 generate all cycles of the theta-graph. There is a unique arm A of the theta-graph containing the two endvertices of the path Q. Let o_3' be the cycle obtained from the path Q by joining its two endvertices in that arm. The cycle o_3' is generated by the cycles o_3, o_1, o_2 and u. To see that note that one of o_3 or $o_3 + u$ is a closed walk that traverses the path P exactly once but does not traverse the path Q. By adding the cycles o_1 and o_2 we can ensure that this closed walk intersects the theta-graph only in the arm A. Thus we have shown that o_3, o_1, o_2, u generate a closed walk that traverses P exactly once and otherwise only uses edges from the arm A; so this closed walk is equal to o_3'.

Let $u' = u + o_3'$, which is a cycle containing the path P but not the path Q.

The cycles o_1, o_2, o_3' and u' clearly generate all cycles of G. Hence o_1, o_2, o_3 and u generate all cycles of the graph G. □

The 'In particular'-part follows from Lemma 2.1. □

Lemma 4.21. *Let G be a graph as in Lemma 4.16.*
Then G contains a graph in the class $\mathcal{W} \cup \mathcal{W}^*$.

Proof that Lemma 4.21 implies Lemma 4.16. This is a direct consequence of Lemma 4.18 and Lemma 4.19. □

We prepare to prove Lemma 4.21. Let θ, o and P be as in Lemma 4.16.

An *arc* is a nontrivial[8] subpath Q of o such that Q intersects the theta-graph θ precisely in its endvertices. We say that Q is a $(\theta\text{-})$*bridge* if its endvertices are interior

[8] A path is *nontrivial* if it contains at least one edge.

vertices of different arms of the theta-graph θ. Otherwise, we say that Q is a (θ-)*detour*. Note that the path P from the assumption includes a bridge, so there is at least one bridge.

Remark 4.22. In this proof we will step by step improve the theta-graph θ and the cycle o. Roughly speaking, this means that we will eliminate the arcs one by one – until at most two of them are left over. Then we will find a configuration in the classes $\mathcal{W}$ or $\mathcal{W}^*$. We start by describing relevant properties of arcs.

No two arcs have adjacent internal vertices (here a vertex of a path is internal if it is not an endvertex). We say that an arc R is *adjacent* to an arc S if there is a subpath X of o joining some of their endvertices such that this subpath does not contain any internal vertices of arcs.

Example 4.23. If there are at least three arcs, then each arc is adjacent to precisely two other arcs. These two arcs are distinct.

We say that an arc R is *weakly adjacent* to an arc S if it is adjacent to S and the endvertex of R on an arc X witnessing adjacency is a branching vertex of θ. It is *strongly adjacent* if the endvertex of R on such a path X is not a branching vertex of θ.

Example 4.24. An adjacent arc that is not weakly adjacent is strongly adjacent. If there are at least three arcs, then no two arcs can be strongly adjacent and weakly adjacent.

The endvertices of a detour Q are contained in a single arm. This arm is uniquely determined – unless the two endvertices of the detour are the two branching vertices of θ. In this case, we choose the unique arm that does not contain any endvertex of the path P. We refer to this uniquely defined arm as the arm *circumvented* by Q (relative to the bridge P). The *replacement path* of Q is the unique subpath of the circumvented arm between the two endvertices of Q. We denote the replacement path of Q by Q'.

Given a bridge B, a detour Q is *B-free* if its replacement path Q' has no internal vertex that is an endvertex of B.

Lemma 4.25. *Assume there is exactly one bridge B, and at least one detour. Then there is a detour that is B-free or that is strongly adjacent to B – or else $\theta \cup o$ is in the class $\mathcal{W}^*$.*

Proof. By assumption there is a detour Q that is adjacent to the bridge B. We may assume, and we do assume, that the detour Q is weakly adjacent and not B-free. In particular, one endvertex of Q is not a branching vertex of θ. Thus if B and Q are the only arcs, then the graph $\theta \cup o$ is in the class $\mathcal{W}^*$. Hence we may assume, and we do assume, that there is another detour. By Example 4.23, there is a detour R adjacent to

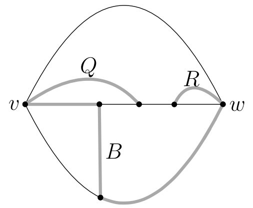

Fig. 6. The path S is highlighted in grey.

B that is distinct from Q. By relabelling the branching vertices if necessary, we may assume, and we do assume, that the branching vertex v is an endvertex of the detour Q.

Next we construct a subtrail S of the cycle o, see Fig. 6 (this subtrail of o is a subpath of o or equal to o). Start the detour Q at the endvertex different from the branching vertex v all the way to the vertex v, then follow the replacement path Q' until you hit an endvertex of B (which must happen eventually as the detour Q is not B-free by assumption), take the bridge B, follow the cycle o until you hit the detour R, then follow R. Denote this trail by S.

We refer to the first vertex of the trail S on the detour R by w. If the vertex w is not a branching vertex, then R is strongly adjacent to B, and we are done. Hence we may assume, and we do assume, that w is a branching vertex. As S is a trail, its interior vertices v and w cannot be identical.

Hence the two endvertices of the trail S cannot be branching vertices, and thus are interior vertices of arms of θ. As the path $o \setminus S$ (between the two endvertices of S) does not include a bridge and no branching vertex of θ, this path must have both its endvertices on a single arm of θ.

Denote the endvertex of the detour R different from w by z. By the above, the vertex z is on the arm of the theta-graph θ circumvented by Q. Note that the replacement path Q' is equal to the subpath of this arm from v to the endvertex of B on that arm. As S is a trail, its endvertex z cannot be on the subpath Q'. So the replacement path for R is disjoint from the path Q', and so in particular does not contain the endvertex of the bridge B on the arm circumvented by R. Thus the detour R is B-free. This completes the proof. $\square$

Lemma 4.26. *If there is at most one bridge, then $\theta \cup o$ has a subgraph in the class $\mathcal{W} \cup \mathcal{W}^*$.*

Proof. From the graph $H = \theta \cup o$ we pick a theta-graph θ' of parameter r and a cycle o' of length at most r containing a branching vertex of θ' so that o' has only a single bridge. Such a choice is possible as we could simply take θ and o. Now we pick o' and θ' amongst all possible choices so that there are as few detours as possible. By replacing '(θ, o)' by '(θ', o')' if necessary, we may assume, and we do assume, that (θ, o) has as few detours as possible. Denote the unique θ-bridge included in o by P. Assume that the graph H has no subgraph in the class $\mathcal{W}^*$.

Suppose for a contradiction, there is a θ-detour included in the cycle o. Let Q be a detour. By Lemma 4.25, we may assume, and we do assume, that Q is P-free or strongly adjacent to P. We denote the arm circumvented by Q by A, and the replacement path for Q by Q'.

We distinguish two cases.

Case 1: the length of Q' is at most the length of Q. We obtain o' from o by replacing the path Q by Q'. As the bridge P and the detour Q are internally disjoint, the closed walk o' includes the path P. We obtain o'' from o' by taking a cycle included in the closed walk o' that includes the path P. This ensures that P is a θ-bridge included in o''. Now let R be a subpath of o'' that intersects θ precisely at its endvertices. As $o'' \setminus \theta \subseteq o$, this subpath R is a subset of the cycle o. Thus R is equal to the bridge P or a detour included in o. We conclude that P is the only θ-bridge included in o'' and there are strictly less θ-detours included in o''. It follows that the path $o'' \setminus P$ includes a branching vertex of θ. Hence the pair (θ, o'') contradicts the minimality of (θ, o). Thus we get a contradiction in this case.

Case 2: the length of Q' is strictly larger than the length of Q. We obtain θ' from the theta-graph θ by replacing the path Q' by Q. As the path Q does not contain any other vertices of θ except for its endvertices, the graph θ' is a theta-graph. And its parameter is at most that of θ. Note that the theta-graphs θ and θ' have the same branching vertices. In particular, the cycle o includes a branching vertex of the theta-graph θ'. By our choice of Q according to Lemma 4.25, we have to consider the following two subcases.

Case 2A: the detour Q is P-free. As the bridge P and the detour Q are internally disjoint, the path P is a θ'-bridge included in o. And it is the only θ'-bridge. All θ'-detours include some θ-detour aside from Q. Thus there are less θ'-detours than θ-detours. This gives a contradiction to the choice of (θ, o) in this subcase.

Case 2B: the detour Q is strongly adjacent to P. We obtain P' from P by adding a subpath X from an endvertex of P to an endvertex of Q witnessing that P and Q are strongly adjacent. Hence the path P' does not contain any branching vertex, and so is a θ'-bridge. All other θ'-arcs are included in o and are disjoint from the subpath X of o, and include θ-arcs. These θ-arcs are θ-detours. So P' is the only θ'-bridge. As Q is not included in a θ'-arc, this gives a contradiction to the choice of (θ, o) in this subcase.

Having considered all cases, we conclude that there is no detour. Hence the graph $\theta \cup o$ is in the class $\mathcal{W}$. This completes the proof. $\square$

We refer to one of the branching vertices included in the cycle o as v. We say that a bridge Q is *primary* if it has an endvertex x such that a shortest path within the cycle o from x to the vertex v includes the bridge Q. A bridge that is not primary is *secondary*.

Lemma 4.27. *If there is a primary bridge Q, then there is a cycle u of length at most r containing v such that Q is the only θ-bridge of u.*

Proof. As the bridge Q is primary, it has an endvertex x such that the shortest path S from x to v within the cycle o includes the primary bridge Q. Take a shortest subpath S' of the path S starting at v that includes a bridge. By minimality, the path S' includes only a single bridge, and edges of S' outside this bridge are on the theta-graph θ or in detours. Denote that single bridge on S' by Q' and the endvertex of S' different from v by x'. As the theta-graph θ has parameter r, it includes a path R from x' to v of length at most $r/2$. The concatenation of the paths S' and R is a closed walk of length at most r that traverses the bridge Q' once. It contains the vertex v. Hence this closed walk includes a cycle u including the bridge Q'. By construction, this cycle u has exactly one bridge. $\square$

Lemma 4.28. *There cannot be two secondary bridges.*

Proof. Either there is a single vertex of the cycle o that has maximum distance from the vertex v in the cycle o, or there is an edge of o such that its two endvertices are the only vertices with maximum distance from v. We refer to these vertices as *barriers*. As different bridges do not have adjacent interior vertices (and if there are two barriers they are adjacent), there can be at most one bridge that has a barrier as an interior vertex.

It suffices to show that any bridge Q that has no barrier as an interior vertex is primary. Let x and y be the two endvertices of the bridge Q. Now consider the path $o - v$. In the middle we have the barriers. As the bridge Q has no barrier as an internal vertex, the vertices x and y must be on the same side of the barriers on $o - v$. Take the vertex of x or y that is nearest to the barriers. The shortest path to v in o from that vertex includes the bridge Q. Thus Q is a primary bridge. Hence all but at most one bridge is primary. $\square$

We conclude this subsection as follows.

Proof of Lemma 4.21. Let G be a graph that contains a theta-graph θ of parameter r such that there is a cycle o of length at most r containing a branching vertex of θ and a path P between interior vertices of different arms of θ that avoids the branching vertices of θ. Let H be the subgraph of G obtained by taking the union of the theta-graph θ and the cycle o.

If there is a primary bridge, then by Lemma 4.27 we can modify the cycle o so that there is only one θ-bridge. Otherwise by Lemma 4.28, there is only a single bridge. In either case, we may assume, and we do assume, that there is only a single bridge. By Lemma 4.26, $\theta \cup o$ has a subgraph in the class $\mathcal{W} \cup \mathcal{W}^*$. $\square$

As we have shown after the statement of Lemma 4.21 above that it implies Lemma 4.16, we have also completed the proof of that lemma.

4.5. Finding bounded wheels

In this subsection we prove two lemmas, which allow us to find bounded wheels in certain situations. They are used in the proof of Lemma 3.7.

Roughly speaking, the next lemma says that we can improve a given theta-graph or else we find a bounded subdivision of a wheel.

Lemma 4.29. *Assume G has a theta-graph θ of parameter r with branching vertices v and w. Let o be a cycle of length at most r including a shortest v-w-path. Then there is a theta-graph θ' of parameter r whose branching vertices are v and w that includes the cycle o – or G includes an r-local subdivision of a wheel.*

Proof. A *weak θ-bridge* is a path P that does not contain any branching vertices of θ but joins interior vertices of different arms of θ. If the cycle o included a weak θ-bridge, then G includes an r-local subdivision of a wheel by Lemma 4.16.

Hence we may assume, and we do assume, that o includes no weak θ-bridge. Hence each of the two v-w-paths included in o intersects at most one arm of θ at interior vertices. So there is an arm Q of θ that intersects the cycle o precisely in the vertices v and w. So $o \cup Q$ is a theta-graph. By assumption, the cycle o includes a shortest v-w-path; call it R. To see that $o \cup Q$ is a theta-graph of parameter r, we show that the cycle RQ has length at most r. Let o' be a cycle including Q of length at most r included in the old theta-graph θ. As R is a shortest v-w-path, the length of RQ is at most that of o'. So RQ has length at most r. Thus $o \cup Q$ is a theta-graph of parameter r, as desired. □

Roughly speaking, the next lemma gives conditions under which a bounded wheel can be built from a bounded fan.

Lemma 4.30. *Assume G has a cycle o of length at most r. Assume there is a fan F of parameter r centered at a vertex v of o whose ending and starting edges are the two edges incident with v on o. Assume there is a vertex w of o not contained in the fan F. Assume no piece of F contains interior vertices of both v-w-paths included in o.*

Then G has an r-local subdivision of a wheel.

Proof. We denote the two v-w-paths included in o by P_1 and P_2. If a piece o_i of the fan F contains an interior vertex of a path P_k, then we colour that piece with that path P_k. By assumption each of piece o_i is coloured with at most one path P_k, while some may not be coloured at all.

As the first and last piece of the fan F are coloured with different paths P_k, the fan F has at least two pieces. Moreover, there are two pieces o_i and o_j with $i < j$ that are coloured with different paths P_k such that all pieces o_m in between (that is, with $i < m < j$) are not coloured at all. Now we consider the subfan $o_i, ..., o_j$ of the original fan F. Denote that subfan by F'.

We claim that the subgraph of G that is the union of the cycle o and the fan F' includes an r-local subdivision of a wheel. The center of this wheel is the vertex v, its pieces are the pieces o_m of the fan F' with $i < m < j$, cycles o_i' and o_j' constructed from o_i and o_j, respectively, and a cycle o' constructed from o as follows.

By symmetry, we assume that the piece o_i contains interior vertices of the path P_1, and the piece o_j contains interior vertices of the path P_2. Let a_1 be a vertex of the piece o_i on the path P_1 that is nearest to the vertex w on P_1; note that this uniquely defines the vertex a_1.

As the path $P_1 - v$ is disjoint from the piece o_{i+1}, the vertex a_1 is an interior vertex of the subpath L_iM_i of the piece o_i. Let Q_1 be the subpath of the path L_iM_i from v to a_1.

Similarly, let a_2 be a vertex of the piece o_j on the path P_2 that is nearest to the vertex w on P_2. And let Q_2 be the subpath of the path M_jR_j from a_2 to v.

Given $k \in \{1, 2\}$, one of the paths vP_ka_k and Q_k is not longer than the other; pick such a path and denote it by S_k.

Sublemma 4.31. *Given $k \in \{1, 2\}$, the path S_k intersects a piece o_m with $i \leq m \leq j$ only in the vertex v – unless $k = 1$ and $m = i$ or else $k = 2$ and $m = j$. The paths S_1 and S_2 intersect only at the vertex v.*

Proof. By symmetry, it suffices to consider the case where $k = 1$. If the path S_1 is chosen to be a subpath of the path P_1 (which is included in the cycle o), the sublemma is immediate. So we may assume, and we do assume that the path S_1 is a subpath of the piece o_i. As a piece o_m with $i < m \leq j$ does not intersect the subpath L_iM_i of the piece o_i in interior vertices, the piece o_m intersects the path S_1 only in the vertex v. It remains to show that the paths S_1 and S_2 intersect only in the vertex v. If one of them is a subpath of a path P_k, this is immediate. Otherwise, the paths S_1 and S_2 are proper subpaths of the paths L_iM_i and M_jR_j containing the vertex v, and thus can only intersect in the vertex v. □

We obtain the cycle o_i' from o_i by replacing the path Q_1 by S_1. Similarly, we obtain the cycle o_j' from o_j by replacing the path Q_2 by S_2. We obtain the cycle o' from o by replacing the paths vP_ka_k by the paths S_k (for $k = 1, 2$).

By the choice of the vertex a_k and Sublemma 4.31, the cycle o' intersects the cycle o_i' precisely in the path S_1, and o' intersects the cycle o_j' precisely in the cycle S_2.

By Sublemma 4.31, the cycles o', o_i', o_j' and o_m with $i < m < j$ are the pieces of a subdivision of a wheel. By construction all these pieces have length at most r. Thus this subdivision of a wheel is r-local. □

4.6. Final step

Proof of Lemma 3.7. Let G be an r-locally 3-connected graph. Our aim is to find an r-local subdivision of a wheel. By Theorem 4.2, the graph G is triangular with parameter

r or has a geodesic cycle of length at most r with at least four edges. If it is triangular, we are done by Lemma 4.5.

Hence we may assume, and we do assume, that G has a geodesic cycle of length at most r with at least four edges. Denote that cycle by o'. Pick two vertices v_0 and v_1 that are not adjacent on the cycle o'. Let P be a shortest v_0-v_1-path included in the geodesic cycle o'. Now apply Lemma 4.10 to o', v_0 and v_1. We get a pre-fan F of parameter r centered at some vertex v_i none of whose pieces contains the vertex v_{i+1} (for some $i \in \mathbb{F}_2$). And there is a geodesic cycle o of length at most r including P such that the start and end of the pre-fan are the neighbours of v_i on o. By Lemma 4.8, we may assume, and we do assume, that the pre-fan F is a fan. By symmetry, we may assume, and we do assume, that the center of the fan F is v_0.

As we would be done otherwise, by Lemma 4.30 there is a piece o_i of the fan F that contains interior vertices of the two v_0-v_1-paths included in the cycle o. Indeed, all other assumptions of that lemma are satisfied by o and F.

Sublemma 4.32. *There is a theta-graph θ of parameter r containing a shortest path between its branching vertices with at least two edges.*

Proof. Let Q be a subpath of the cycle o_i that intersects the cycle o precisely at its endvertices such that these endvertices are on different v_0-v_1-paths included in the cycle o.

The graph $o \cup Q$ is a theta-graph. As the cycle o is geodesic, it includes a shortest path between the branching vertices. The cycle obtained by concatenating this path with the path Q cannot be longer than the cycle o_i. Thus the theta-graph $o \cup Q$ has parameter r.

As the cycle o is geodesic, it contains a shortest path between the two branching vertices. Each of the two paths between the branching vertices included in the cycle o contains one of the vertices v_0 or v_1. Hence the theta-graph $o \cup Q$ includes a shortest path with at least two edges. □

Let θ be a theta-graph as in Sublemma 4.32. Denote its two branching vertices by $\bar{v}$ and $\bar{w}$ and a shortest path between them by $\bar{P}$. Let $\bar{o}$ be a cycle of θ including $\bar{P}$ that does not include an edge between the branching vertices $\bar{v}$ and $\bar{w}$ (such a choice is possible as θ contains two paths between the vertices $\bar{v}$ and $\bar{w}$ aside from $\bar{P}$). Now we apply Lemma 4.10 to the cycle $\bar{o}$, the path $\bar{P}$ and the vertices $\bar{v}$ and $\bar{w}$. This part of the proof is similar to the application above but now with the theta-graph θ, we have slightly stronger assumptions and will hence be able to complete the proof.

We get a pre-fan $\bar{F}$ of parameter r centered at one of the vertices $\bar{v}$ or $\bar{w}$, say $\bar{v}$, none of whose pieces contains the vertex $\bar{w}$. And there is a cycle $\bar{u}$ of length at most r including $\bar{P}$ such that the start and end of the pre-fan are the neighbours of $\bar{v}$ on $\bar{u}$. By Lemma 4.8, we may assume, and we do assume, that the pre-fan $\bar{F}$ is a fan.

As we would be done otherwise, by Lemma 4.30 there is a piece $\bar{o}_i$ of the fan $\bar{F}$ that contains interior vertices of the two $\bar{v}$-$\bar{w}$-paths included in the cycle $\bar{u}$. By replacing the theta-graph θ by the theta-graph guaranteed by Lemma 4.29 if necessary, we may assume, and we do assume, that the cycle $\bar{u}$ is included in the theta-graph θ. By Lemma 4.16 applied to θ and a suitable subpath of the piece $\bar{o}_i$, the graph G includes an r-local subdivision of a wheel. $\square$

5. Concluding remarks

Splitter theorems have turned out to be a key tool when it comes to characterising certain minor closed classes of graphs or matroids. And in fact they can be proven in fairly general settings. For example Chun, Mayhew and Oxley proved a splitter theorem for internally 4-connected binary matroids [4]. We expect that there is the following splitter theorem for r-locally 3-connected graphs.

Conjecture 5.1. *Every r-locally 3-connected graphs contains an edge to delete or contract such that r-local 3-connectivity is preserved — unless G is K_4.*

A *cycle-decomposition* is a graph-decomposition whose decomposition graph is a cycle. For example, every path-decomposition is a cycle decomposition. A natural future application of Theorem 1.1 might be to characterise (r-locally 2-connected) graphs that have a cycle decomposition of width at most two and locality at least r. It is expected that the characterisation of graphs of path-width at most two [9], see also [2], extends to this setting in the natural way.

More generally, it is expected that Theorem 1.1 can be used to extend excluded minors characterisations for any minor-closed subclass of the class of series-parallel graphs to our new setting of local separators.

Another direction in which one could try to extend Theorem 1.1 is to characterise graphs with graph-decompositions of width at most three and locality at least r. The starting point would be the characterisation of graphs of tree-width at most three by Arnborg, Proskurowski & Corneil (1990) and Satyanarayana & Tung (1990) [1,15].

Appendix A. Basic definitions

In this appendix, we recall some notions from [3].

There are various equivalent definitions of tree-decompositions:

1. the one in terms of a tree displaying the global structure of the graph by encoding how certain bags are glued onto one another;
2. a recursive definition in terms of nested gluing; and
3. a decomposition perspective, where in one decomposition step one separates the graph along a separator and replaces it by a complete graph.

All these notions of tree-decompositions naturally generalise to r-local graph-decompositions. In this paper it is easiest to work with the decomposition perspective. So we now start setting up the terminology to do that.

Given a graph G with a vertex v and an integer s, the *ball* of radius s around the vertex v is the induced subgraph of G, whose vertices are those of distance at most s from v and without all edges joining two vertices of distance precisely s. Similarly, given a half-integer $s+\frac{1}{2}$, the *ball* of radius $s+\frac{1}{2}$ around the vertex v is the induced subgraph of G, whose vertices are those of distance at most s from v. We denote the ball of radius s around v by $B_s(v)$. Below we will often consider the graph $B_s(v)-v$, to which we refer as a *punctured ball.* Given a parameter $r \in \mathbb{N}\cup\{\infty\}$, a vertex v is an *r-local cutvertex* if it separates the ball of radius $r/2$ around v; formally: $B_{r/2}(v)-v$ is disconnected.

Given a parameter $r \geq 1$ and a graph G with a vertex v, the graph obtained from G by *r-locally cutting* the vertex v is defined as follows. Let X be the set of connected components of the ball of radius r around v with v removed; formally X is the set of components of the graph $B_{r/2}(v)-v$. Define a new graph from G by replacing the vertex v by one new vertex for each element of the set X, where the vertex labelled with $x \in X$ inherits the incidences with those edges incident with v that are incident with a vertex of the connected component X. We refer to the new vertices as the *slices of v*. This completes the construction of the r-local cutting of G.

Given a graph G and a parameter r, we say that G has an *r-local graph-decomposition of adhesion one* with set of bags $\mathcal{B}$ if we can iteratively apply r-local cuttings to G such that the resulting graph is a disjoint union of the family $\mathcal{B}$. Recall the definitions of local 2-connectedness from the first paragraphs of Section 2. In [3], we prove the following local refinement of the block-cutvertex theorem:

Theorem A.1 *([3, Theorem B.1]). Every r-locally connected graph has an r-local graph-decomposition of adhesion one such that all its bags are r-locally 2-connected.*

Now we give a formal definition of the *explorer-neighbourhood* of parameter r with explorers based at the vertices v and w with distance[9] at most $\frac{r}{2}$. The core is the set of all vertices on shortest paths between the vertices v and w. We take a copy of the ball $B_{r/2}(v)$ where we label a vertex u with the set of shortest paths from the core to u contained in the ball $B_{r/2}(v)$. Similarly, we take a copy of the ball $B_{r/2}(w)$ where we label a vertex u with the set of shortest paths from the core to u contained in the ball $B_{r/2}(w)$. Now we take the union of these two labelled balls – with the convention that two vertices are identified if they have a common label in their sets; that is, there is a shortest path from the core to that vertex discovered by both explorers. (Note that the same vertex x of G could be in both balls but the label sets could be disjoint, see Fig. 4. In this case there would be two copies of that vertex in the union. In

[9] In this paper the explorer-neighbourhood of vertices v and w of distance more than $\frac{r}{2}$ is undefined; and hence throughout the paper in statements where the explorer-neighbourhood is mentioned we have implicitly the assumption that the involved vertices have distance at most $\frac{r}{2}$.

such a case the union would not be a subgraph of the original graph.) We denote the explorer neighbourhood by $\mathrm{Ex}_{\mathrm{r}}(v, w)$. This completes the definition of explorer neighbourhood.

A set $\{v, w\}$ is an *r-local 2-separator* if the punctured explorer-neighbourhood $\mathrm{Ex}_{\mathrm{r}}(v, w) - v - w$ is disconnected; and the vertices v and w have distance at most $r/2$. Given a graph G with an r-local 2-separator $\{v_0, v_1\}$, the graph obtained from G by *r-locally cutting* $\{v_0, v_1\}$ is defined as follows. Let X be the set of connected components of the punctured explorer-neighbourhood $\mathrm{Ex}_{\mathrm{r}}(v_0, v_1) - v_0 - v_1$. We now replace in the graph G the vertices v_0 and v_1 each by one copy for every element of X. Here a copy of v_i labelled by some $x \in X$ inherits an edge from v_i if the other endvertex of that edge is a vertex of the component x. We refer to the newly added vertices as the *slices* of the vertices v_1 or v_2, respectively. We additionally add a weighted edge between any two slices for the same $x \in X$. Its weight is given by the minimum length of a path between v_0 and v_1 in the explorer-neighbourhood $\mathrm{Ex}_{\mathrm{r}}(v_0, v_1)$ with the component x removed. It follows that all but one of these weights are always the same. We refer to these additional edges as *torso edges*. If the vertices v_0 and v_1 share an edge e in G, we add a new connected component consisting of the edge e and one edge in parallel to e. This other edge is a torso edge and its length is the minimum length of a path between v_0 and v_1 in the explorer-neighbourhood $\mathrm{Ex}_{\mathrm{r}}(v_0, v_1)$. Finally, we replace each torso edge by a path of the same length; we refer to such paths as *torso-paths*.[10] This completes the definition of local cutting.

Given an r-locally 2-connected graph G and a parameter r, we say that G has an *r-local graph-decomposition of adhesion two* with set of bags $\mathcal{B}$ if we can iteratively apply r-local cuttings at r-local 2-separators to G such that the resulting graph is a disjoint union of the family $\mathcal{B}$. Recall the definitions of local connectedness from the first two paragraphs of Section 2. In [3], we prove the following local extension of the decomposition theorem along 2-separators:

Theorem A.2 *([3, Theorem 1.2]). Every r-locally 2-connected graph has an r-local graph-decomposition of adhesion two such that all its torsos are r-locally 3-connected or cycles of length at most r.*

We say that a graph-decomposition with family of bags $\mathcal{B}$ has width at most w if all members of $\mathcal{B}$ have size at most $w - 1$.

Finally, we summarise some basic notation concerning generation by families of cycles. A *walk* is a sequence alternating between edges and vertices such that adjacent pairs are incident and such that they end and start with a vertex; these vertices are called the *start vertex* and *end vertex*, respectively. Note that in a walk edges and vertices may appear multiple times. Two walks such that the endvertex of the first walk is equal to the start

[10] This technical step reduces technicalities elsewhere; indeed, the explorer-neighbourhood is not defined for weighted graphs, and doing so would lead to technicalities.

vertex of the second walk can be *concatenated* by gluing together their two sequences by identifying the end vertex of the first walk with the start vertex of the second one. Given vertices v and w, an edge e and a walk W that has the subwalk $wevew$, the *reduction* of W by $wevew$ is the sequence obtained from W by replacing the walk $wevew$ by the trivial walk w. Clearly, reductions are walks with the same start vertices and end vertices.

A walk is *closed* if its start vertex and end vertex are the same. Sometimes closed walks are also defined through the associated cyclic sequences obtained from the linear sequence by merging the first and last vertex; in this paper we implicitly move between these two representations. For example, a *cycle* is a closed walk where each edges appears at most once and each vertex appears at most once (here we refer to the second representation of closed walks, of course). We say that a family of closed walks $\mathcal{W}$ *generates* a closed walk X (homotopically) if there is a concatenation of members of $\mathcal{W}$ that can be reduced to X.

This topological definition of generating is complemented by the following algebraic variant. Given a graph with edge set E and a field k, consider the vector space k^E, whose vectors have components that are indexed by the edges of the graph. Given a set of cycles $\mathcal{C}$ and a field k, we say that $\mathcal{C}$ *generates a cycle X over the field k* if the family of characteristic vectors of elements of $\mathcal{C}$ generates the characteristic vector of X in the vector space k^E. It is straightforward to check that if a family of cycles generates a cycle homotopically, then it generates it over any field, in particular the two-element field $\mathbb{F}_2$.

References

[1] Stefan Arnborg, Andrzej Proskurowski, Derek G. Corneil, Forbidden minors characterization of partial 3-trees, Discrete Math. 80 (1) (1990) 1–19.

[2] János Barát, Peter Hajnal, Yixun Lin, Aifeng Yang, On the structure of graphs with path-width at most two, Studia Sci. Math. Hung. 49 (2009) 10.

[3] Johannes Carmesin, Local 2-separators, preprint, available at http://web.mat.bham.ac.uk/J.Carmesin/.

[4] Carolyn Chun, Dillon Mayhew, James Oxley, Towards a splitter theorem for internally 4-connected binary matroids, J. Comb. Theory, Ser. B 102 (3) (2012) 688–700.

[5] Julia Chuzhoy, Zihan Tan, Towards tight(er) bounds for the excluded grid theorem, J. Comb. Theory, Ser. B 146 (2021) 219–265.

[6] R. Diestel, K.Yu. Gorbunov, T. Jensen, C. Thomassen, Highly connected sets and the excluded grid theorem, J. Comb. Theory, Ser. B 75 (1999) 61–73.

[7] David Eppstein, Grid minors in damaged grids, Electron. J. Comb. 21 (3) (2014).

[8] Jakub Gajarský, Stephan Kreutzer, Jaroslav Nešetřil, Patrice Ossona de Mendez, Michał Pilipczuk, Sebastian Siebertz, Szymon Toruńczyk, First-order interpretations of bounded expansion classes, ACM Trans. Comput. Log. 21 (4) (July 2020).

[9] Nancy G. Kinnersley, Michael A. Langston, Obstruction set isolation for the gate matrix layout problem, Discrete Appl. Math. 54 (2–3) (1994) 169–213.

[10] Matthias Kriesell, Unique colorability and clique minors, J. Graph Theory 85 (1) (2017) 207–216.

[11] Irene Muzi, Michael P. O'Brien, Felix Reidl, Blair D. Sullivan, Being even slightly shallow makes life hard, in: MFCS, vol. 83, 2017, pp. 79:1–79:13.

[12] Jaroslav Nešetřil, Patrice Ossona de Mendez, Sparsity: Graphs, Structures, and Algorithms, Springer, 2012.

[13] Serge A. Plotkin, Satish Rao, Warren D. Smith, Shallow excluded minors and improved graph decompositions, in: SODA, vol. 90, 1994, pp. 462–470.

[14] Neil Robertson, Paul D. Seymour, Graph minors. v. Excluding a planar graph, J. Comb. Theory, Ser. B 41 (1) (1986) 92–114.
[15] A. Satyanarayana, L. Tung, A characterization of partial 3-trees, Networks 20 (3) (1990) 299–322.
[16] Carsten Thomassen, A minimal condition implying a special K_4-subdivision in a graph, Arch. Math. 25 (1) (1974) 210–215.