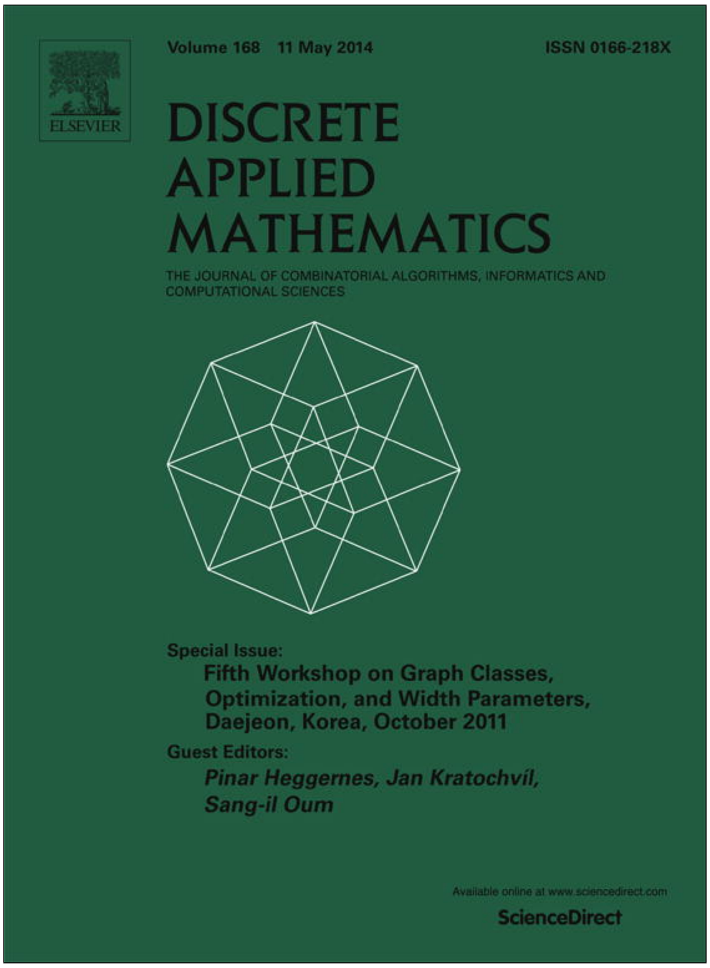
Volume 168 11 May 2014
ISSN 0166-218X
ELSEVIER
DISCRETE
APPLIED
MATHEMATICS
THE JOURNAL OF COMBINATORIAL ALGORITHMS, INFORMATICS AND COMPUTATIONAL SCIENCES
Special Issue:
Fifth Workshop on Graph Classes, Optimization, and Width Parameters, Daejeon, Korea, October 2011
Guest Editors:
Pinar Heggernes, Jan Kratochvíl, Sang-il Oum
Available online at www.sciencedirect.com
ScienceDirect

Discrete Applied Mathematics 168 (2014) 34–39

Contents lists available at ScienceDirect

Discrete Applied Mathematics

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

Square roots of minor closed graph classes

Nestor V. Nestoridis, Dimitrios M. Thilikos [*,1]

Department of Mathematics, National and Kapodistrian University of Athens, Panepistimioupolis, GR-15784 Athens, Greece

ARTICLE INFO

Article history:
Received 29 February 2012
Received in revised form 23 May 2013
Accepted 27 May 2013
Available online 22 June 2013

Keywords:
Square roots of graphs
Branch-width
Carving-width
Graph minors

ABSTRACT

Let $\mathcal{G}$ be a graph class. The *square root* of $\mathcal{G}$ contains all graphs whose squares belong in $\mathcal{G}$. We prove that if $\mathcal{G}$ is non-trivial and minor closed, then all graphs in its square root have carving-width bounded by some constant depending only on $\mathcal{G}$. As a consequence, every square root of such a graph class has a linear time recognition algorithm.

1. Introduction

Let $\mathcal{G}$ be a graph class. The *square root* of $\mathcal{G}$ is defined as the graph class

$$\sqrt{\mathcal{G}} = \{G \mid G^2 \in \mathcal{G}\},$$

where the *square* G^2 of a graph G is the graph obtained from G after adding edges between all pairs of vertices that share a common neighbor.

In [3], Harary, Karp, and Tutte provided a complete characterization of the graphs in $\sqrt{\mathcal{P}}$ where $\mathcal{P}$ is the class of all planar graphs. Notice that planar graphs are *minor closed*, i.e. a minor of every graph in $\mathcal{P}$ also belongs in $\mathcal{P}$. Minor closeness is a very general property that is satisfied by a great variety of graph classes; see e.g. [6].

According to the characterization of [3], all graphs in $\sqrt{\mathcal{P}}$ are outerplanar and have bounded degree. This implies that graphs in $\sqrt{\mathcal{P}}$ have a very specific "tree-like" structure and it is a natural question whether this is the case for more general minor closed graph classes. In this paper we extend this result, in the sense that the same tree-like property holds for every minor closed graph class that is non-trivial (i.e. that does not contain all graphs). In fact, we prove (in Section 3) that, in this case, the correct "tree-likeness" property is given by the parameter of carving-width, introduced by Seymour and Thomas in [11]. As a consequence, we prove in Section 4 that the square root of any non-trivial minor closed graph class has a linear time recognition algorithm. This extends the algorithmic results of [5] where a linear time algorithm was given for recognizing the square roots of planar graphs.

* Corresponding author.
E-mail addresses: nestornestoridis@yahoo.gr (N.V. Nestoridis), sedthilk@thilikos.info (D.M. Thilikos).

1 The second author was Co-financed by the European Union (European Social Fund — ESF) and Greek national funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) - Research Funding Program: "Thales. Investing in knowledge society through the European Social Fund".

0166-218X/$ – see front matter
http://dx.doi.org/10.1016/j.dam.2013.05.026

2. Definitions

We next give some definitions that are necessary in order to formally define carving width. This will permit us to give the formal statement of our combinatorial result.

Boundaries in graphs and hypergraphs. In this paper we deal with graphs and hypergraphs. For a (hyper)graph G we denote by $V(G)$ its vertex set and by $E(G)$ the set of its (hyper)edges. If $S \subseteq V(G)$ (resp. $F \subseteq E(G)$) we denote $\overline{S} = V(G) \setminus S$ (resp. $\overline{F} = E(G) \setminus F$).

Given a vertex set $S \subseteq V(G)$, let $E_G(S)$ be the set of hyperedges containing vertices in S. For simplicity we also denote $E_G(v) = E_G(\{v\})$. We also define $\Delta(G) = \max\{|E_G(v)| \mid v \in V(G)\}$. Given a set $S \subseteq V(G)$, we define

$$\partial_G(S) = E_G(S) \cap E_G(\overline{S}).$$

Notice that ∂_G is a symmetric function, i.e. for every $S \subseteq V(G)$, it holds that $\partial_G(S) = \partial_G(\overline{S})$. Also given a set $F \subseteq E(G)$, we set

$$\partial^*_G(F) = \left(\bigcup_{f \in F} f\right) \cap \left(\bigcup_{f \in \overline{F}} f\right).$$

Given a hypergraph G we define its *dual* as the hypergraph

$$G^* = (E(G), \{E_G(v) \mid v \in V(G)\}).$$

Notice that the hypergraphs G and G^* have the same incidence graph with the roles of their two parts reversed. Given a set $S \subseteq V(G)$ (resp. $F \subseteq E(G)$) we denote by S^* (resp. F^*) their dual hyperedges (resp. vertices) in G^*.

Using duality, we also define $\Delta^*(G) = \Delta(G^*)$. Clearly, for a simple graph G, $\Delta^*(G) = 2$. Moreover, the above definitions imply that for every $F \subseteq E(G)$, $(\partial^*_G(F))^* = \partial_{G^*}(F^*)$.

A graph H is a *minor* of a graph G, and we write $H \leq G$, if H can be obtained for some subgraph of G after contracting edges (the contracting an edge $e = \{x, y\}$ is the operation that removes x and y from G and introduces a new vertex v_e that is made adjacent with all the neighbors of x and y in G, except from x and y). A graph class $\mathcal{G}$ is *minor closed* if every minor of a graph in $\mathcal{G}$ is also a graph in $\mathcal{G}$.

Carving-width. Given a tree T we denote the set of its leaves by $L(T)$ and we call it *ternary* if all vertices in $V(T) \setminus L(T)$ have degree 3. A *carving decomposition* of a hypergraph G is a pair (T, ρ), where T is a ternary tree and ρ is a bijection from $V(G)$ to $L(T)$. The *bridge function* $\beta : E(T) \to 2^{E(G)}$ of a carving decomposition maps every edge e of T to the set $\partial_G(\rho^{-1}(L(T')))$ where T' is one of the two connected components of $T \setminus e$. The *width* of (T, ρ) is equal to $\max_{e \in E(T)} |\beta(e)|$ and the *carving-width* of G, $\mathbf{cw}(G)$, is the minimum width over all carving decompositions of G. The following observation is a direct consequence of the definitions.

Observation 1. *For every hypergraph G, it holds that $\Delta(G) \leq \mathbf{cw}(G)$.*

The main combinatorial result of this paper is the following.

Theorem 1. *For every non-trivial minor closed graph class $\mathcal{G}$ there is a constant $c_{\mathcal{G}}$ such that all graphs in $\sqrt{\mathcal{G}}$ have carving-width at most $c_{\mathcal{G}}$.*

The proof of Theorem 1 uses the parameter of branch-width defined in [8].

Branch-width. A *branch decomposition* of a graph G is a pair (T, τ), where T is a ternary tree and τ is a bijection from $E(G)$ to $L(T)$. The *boundary function* $\omega : E(T) \to 2^{V(G)}$ of a branch decomposition maps every edge e of T to the set $\partial^*_G(\rho^{-1}(L(T')))$ where T' is one of the two connected components of $T - \{e\}$. The *width* of (T, τ) is equal to $\max_{e \in E(T)} |\omega(e)|$ and the *branch-width* of G, $\mathbf{bw}(G)$, is the minimum width over all branch decompositions of G.

The following observation is a direct consequence of the duality between the functions ∂_G and ∂^*_G.

Observation 2. *For every hypergraph G it holds that $\mathbf{bw}(G) = \mathbf{cw}(G^*)$.*

3. Walls and squares

Walls. A *wall of height k*, $k \geq 1$, is obtained from a $((k+1) \times (2k+2))$-grid with vertices (x, y), $x \in \{0, \ldots, 2k+1\}$, $y \in \{0, \ldots, k\}$, after removing the "vertical" edges $\{(x, y), (x, y+1)\}$ for odd $x + y$. We denote such a wall by W_k. A *subdivided wall of height k* is obtained by the wall W_k with some edges of W_k replaced by paths without common internal vertices. If, in such a subdivided wall, all edges have been subdivided at least once, then we say that it is *properly subdivided*. We also say that a graph H is *topological minor* of a graph G if some subdivision of H is a subgraph of G.

The following result follows from the results in [7,8].

Proposition 1 ([7,8])**.** *There is a function $g : \mathbb{N} \to \mathbb{N}$ such that every graph G with branchwidth at least $g(k)$ contains the $(k \times k)$-grid as a minor.*

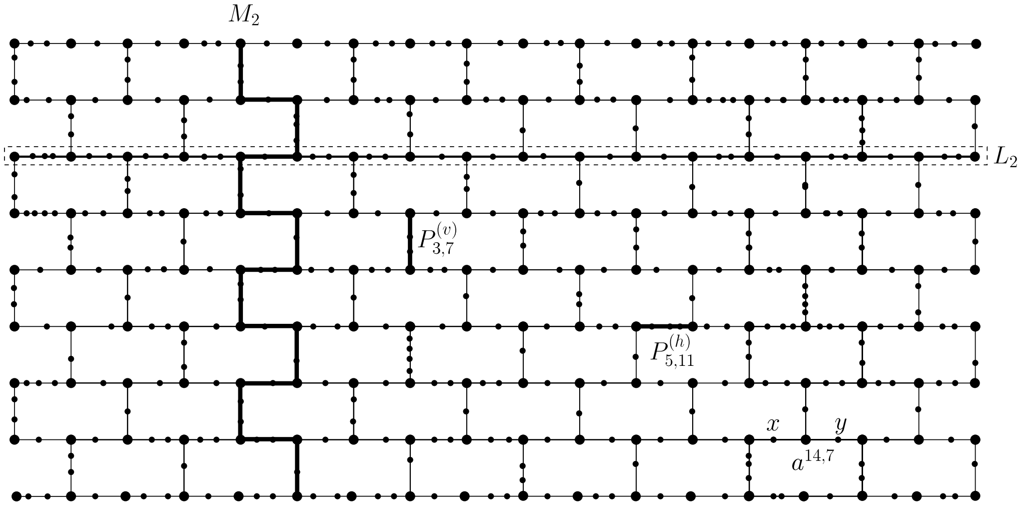

Fig. 1. A properly subdivided wall G of height 8. The big (resp. small) vertices are the original (subdivision) vertices. The paths $P^{(v)}_{3,7}$, $P^{(h)}_{5,11}$, L_2 and the 2-meander M_2 are depicted in G. Notice that $x = \mathbf{left}_G(a^{14,7})$ and $y = \mathbf{right}_G(a^{14,7})$.

A direct consequence of Proposition 1 is the following.

Lemma 1. *There exists a function f such that every graph G with branchwidth at least $f(k)$ contains a properly subdivided wall of height k as a subgraph.*

Proof. Let W'_k be the graph taken from a k-wall if we subdivide each edge once. Notice that W'_k is a subgraph of the $((4k+3)\times(4k+3))$-grid. The lemma follows from Proposition 1 and the fact that minor relation and topological minor relation are identical when the host graph has maximum degree 3 (notice that $\Delta(W'_k) \leq 3$). □

Let H be a $((k+1)\times(2k+2))$-grid where k is a positive odd integer. Let also W_k be a spanning subgraph that is a k-wall as explained above. Let also G be a subdivision of W_k where the non-subdividing vertices are denoted as coordinates of H (we call these vertices *original*). In particular, we denote the vertex of H with coordinates i, j by $a^{i,j}$ (see Fig. 1). For $i \in \{0, \ldots, k-1\}$ and $j \in \{0, \ldots, 2k+1\}$ where $i + j$ is even, we denote by $P^{(v)}_{i,j}$ the path that has replaced the edge $\{a^{i,j}, a^{i+1,j}\}$ in W_k and we call these paths *vertical*. For $i \in \{0, \ldots, k\}$ and $j \in \{0, \ldots, 2k\}$, we denote by $P^{(h)}_{i,j}$ the path that has replaced the edge $\{a^{i,j}, a^{i,j+1}\}$ and we call these paths *horizontal*. We denote by L_i, $i \in \{0, \ldots, k\}$ the path $P^h_{i,0} \oplus \cdots \oplus P^h_{i,2k}$, we direct it from $a^{i,0}$ to $a^{i,2k}$ and we call it *i-th line* of G. We also define, for $l \in \{0, \ldots, k\}$, the *l-th meander* of G as the path

$$M_l = P^{(v)}_{0,2l} \oplus P^{(h)}_{1,2l} \oplus P^{(v)}_{1,2l+1} \oplus P^{(h)}_{2,2l} \oplus P^{(v)}_{2,2l} \oplus \cdots \oplus P^{(v)}_{k-2,2l+1} \oplus P^{(h)}_{k-1,2l} \oplus P^{(v)}_{k-1,2l}.$$

Finally, for every original vertex $a^{i,j}$ where $0 \leq j \leq 2k$ ($1 \leq j \leq 2k+1$), we define $\mathbf{right}_G(a^{i,j})$ ($\mathbf{left}_G(a^{i,j})$) as the vertex of the directed path L_i that appears right after (before) $a^{i,j}$.

We denote by K_r the complete graph on r vertices. Also we use the notation $K_{r,q}$ for the complete bipartite graph with parts of size r and q.

Lemma 2. *Let k be a positive odd integer. If G is a properly subdivided wall of height k, then $K_{k+2} \leq G^2$.*

Proof. For each $i \in \{0, \ldots, k\}$ we define the path P_i of G^2 as follows: If $L_i = (x_0, x_1, \ldots, x_m)$ is the i-th horizontal path of G, then $\bar{L}_i = (x_0, x_2, x_4, \ldots, x_{m'})$ where $m' = m$ or $m' = m - 1$ depending on whether m is even or not. We set $S = \bigcup_{i\in\{0,\ldots,k\}} V(\bar{L}_i)$.

Our next step is to define the paths Q_j of G^2, $l \in \{0, \ldots, k\}$. Let M_l be the l-th meander of G. Following the definition of a meander, we denote

- $P^{(v)}_{i,2l} = (y^{i,2l}_1, \ldots, y^{i,2l}_{w_{i,2l}})$, for even $i \in \{0, \ldots, k-1\}$ (assuming that $y^{i,2l}_1 = a^{i,2l}$ and $y^{i,2l}_{w_{i,2l}} = a^{i+1,2l}$),
- $P^{(v)}_{i,2l+1} = (y^{i,2l+1}_1, \ldots, y^{i,2l+1}_{w_{i,2l+1}})$, for odd $i \in \{0, \ldots, k-2\}$ (assuming that $y^{i,2l+1}_1 = a^{i,2l+1}$ and $y^{i,2l+1}_{w_{i,2l+1}} = a^{i+1,2l+1}$),
- $P^{(h)}_{i,2l} = (z^{i,2l}_1, \ldots, z^{i,2l}_{m_{i,2l}})$, $i \in \{1, \ldots, k-1\}$ (assuming that $z^{i,2l}_1 = a^{i,2l}$ and $z^{i,2l}_{m_{i,2l}} = a^{i,2l+1}$).

For each even $i \in \{0, \ldots, k-1\}$, we set $\bar{P}^{(v)}_{i,2l} = (\bar{y}^{i,2l}_1, y^{i,2l}_2, \ldots, y^{i,2l}_{w_{i,2l}-1}, \bar{y}^{i,2l}_{w_{i,2l}})$ where

$$\bar{y}^{i,2l}_1 = \begin{cases} y^{i,2l}_1 & \text{if } y^{i,2l}_1 \notin S \\ \mathbf{right}_G(y^{i,2l}_1) & \text{if } y^{i,2l}_1 \in S \end{cases}$$

$$\bar{y}^{i,2l}_{w_{i,2l}} = \begin{cases} y^{i,2l}_{w_{i,2l}} & \text{if } y^{i,2l}_{w_{i,2l}} \notin S \\ \mathbf{right}_G(y^{i,2l}_{w_{i,2l}}) & \text{if } y^{i,2l}_1 \in S, \end{cases}$$

and for each odd $i \in \{0, \ldots, k-1\}$, we set $\bar{P}^{(v)}_{i,2l+1} = (\bar{y}_1^{i,2l+1}, y_2^{i,2l+1}, \ldots, y^{i,2l+1}_{w_{xi,2l}-1}, \bar{y}^{i,2l+1}_{w_{i,2l+1}})$ where

$$\bar{y}_1^{i,2l+1} = \begin{cases} y_1^{i,2l+1} & \text{if } y_1^{i,2l+1} \notin S \\ \mathbf{left}_G(y_1^{i,2l+1}) & \text{if } y_1^{i,2l+1} \in S \end{cases}$$

$$\bar{y}^{i,2l+1}_{w_{i,2l+1}} = \begin{cases} y^{i,2l+1}_{w_{i,2l+1}} & \text{if } y^{i,2l+1}_{w_{i,2l+1}} \notin S \\ \mathbf{left}_G(y^{i,2l+1}_{w_{i,2l+1}}) & \text{if } y^{i,2l+1}_{w_{i,2l+1}} \in S. \end{cases}$$

Observe that for each $i \in \{1, \ldots, k-1\}$ and $l \in \{0, \ldots, k\}$ the vertices in $P^{(h)}_{i,2l} \setminus S$ induce a path in G^2 which we denote by $\bar{P}^{(h)}_{i,2l}$.

Notice, that for every odd $i \in \{1, \ldots, k-2\}$ it holds that

$$V(\bar{P}^{(v)}_{i-1,2l}) \cap V(\bar{P}^{(h)}_{i,2l}) = \{\bar{y}^{i-1,2l}_{w_{i-1,2l}}\} \tag{1}$$

$$V(\bar{P}^{(h)}_{i,2l}) \cap V(P^{(v)}_{i,2l+1}) = \{\bar{y}_1^{i,2l+1}\}. \tag{2}$$

Also, for every even $i \in \{2, \ldots, k-1\}$ it holds that

$$V(\bar{P}^{(v)}_{i-1,2l+1}) \cap V(\bar{P}^{(h)}_{i,2l}) = \{\bar{y}^{i-1,2l+1}_{w_{i-1,2l+1}}\} \tag{3}$$

$$V(\bar{P}^{(h)}_{i,2l}) \cap V(\bar{P}^{(v)}_{i,2l}) = \{\bar{y}_1^{i,2l}\}. \tag{4}$$

From (1)–(4) we obtain that

$$\bar{M}_j = \bar{P}^{(v)}_{0,2l} \oplus \bar{P}^{(h)}_{1,2l} \oplus \bar{P}^{(v)}_{1,2l+1} \oplus \bar{P}^{(h)}_{2,2l} \oplus \bar{P}^{(v)}_{2,2l} \oplus \cdots \oplus \bar{P}^{(v)}_{k-2,2l+1} \oplus \bar{P}^{(h)}_{k-1,2l} \oplus \bar{P}^{(v)}_{k-1,2l}$$

is a path in G^2 for each $l \in \{0, \ldots, k\}$.

As paths in $\bar{\mathcal{M}} = \{\bar{M}_l \mid l \in \{0, \ldots, k\}\}$ avoid all vertices of S, we have that for each $i, l \in \{0, \ldots, k\} V(\bar{L}_i) \cap V(\bar{M}_l) = \emptyset$. Moreover, by construction, any two paths in $\mathcal{M}$ are vertex disjoint. Recall that the same holds for the paths in $\bar{\mathcal{L}} = \{\bar{L}_i \mid i \in \{0, \ldots, k-1\}\}$.

We claim that for each $i, l \in \{0, \ldots, k\}$ there is an edge $e_{i,l}$ in G^2 with one endpoint in $\bar{L}_i$ and the other in $\bar{M}_l$. Indeed it is easy to see that one can take

$$e_{i,l} = \begin{cases} \{a^{i,2l}, \mathbf{left}(a^{i,2l})\} & \text{if } (i, l) \in \{(0, k), (k, k)\} \\ \{a^{i,2l}, \mathbf{right}(a^{i,2l})\} & \text{otherwise.} \end{cases}$$

Now we are ready to prove that G^2 contains K_{k+2} as a minor. For this, we first remove from G^2 all edges that are neither edges of paths in $\bar{\mathcal{M}} \cup \bar{\mathcal{L}}$ nor edges in $\mathcal{E} = \{e_{i,l} \mid (i, l) \in \{0, \ldots, k\}^2\}$. Then we contract all edges of the remaining graph except from those in $\mathcal{E}$ and we end up to a graph isomorphic to $K_{k+1,k+1}$ that, in turn, can be further contracted to K_{k+2}. □

Extended carving-decompositions. Let us relax the definition of a carving decomposition in the case of graphs in the following way: let G be a graph, T be a ternary tree rooted on some, say r, of its leaves, and a surjection $\sigma : L(T) \setminus \{r\} \to V(G)$. As T is rooted each edge of T can be seen as a directed edge pointing towards the root. Let $f = (x, y)$ be such an edge and let T_x and T_y be the two connected components of $T \setminus f$ such that $r \in V(T_y)$. Let also $V_x = \sigma(L(T_x))$ and $V_y = \sigma(L(T_y))$. We define the *bridge set* of f as $\beta(f) = \{\{a, b\} \in E(G) \mid a \in V_x \text{ and } b \in V_y \setminus V_x\}$. We call each such pair (T, σ) an *extended carving decomposition* of G and we define its *width* as the maximum $|\beta(f)|$ over all $f \in E(T)$.

Lemma 3. *If G has an extended carving decomposition of width at most k, then it also has a carving decomposition of width at most k.*

Proof. Let (T, σ) be such an extended carving decomposition. We apply on (T, σ) the following transformation until this is not possible any more: if for some v, in $\sigma^{-1}(v)$ contains at least two vertices a and b then replace σ by $\sigma \setminus \{(a, v)\}$ and replace T by the tree that is obtained if we remove a and then we dissolve the resulting vertex of degree 2. It is easy to see that each new pair (T, σ) has width at most k. Also when the above procedure finishes, σ^{-1} is a bijection from $V(G)$ to $L(T) \setminus \{r\}$. If now we remove r from T and dissolve the resulting vertex of degree 2, we have a branch decomposition (T', σ^{-1}) of G with width at most k. □

Lemma 4. *For every graph G, $\mathbf{cw}(G) \leq \Delta(G) \cdot \mathbf{bw}(G)$.*

Proof. We set $\delta = \Delta(G)$. Let (T, τ) be a branch decomposition of G with width at most k. We now define an extended carving decomposition (T', σ) of G as follows: Take T and for each $v \in L(T)$ we define two new vertices v_1 and v_2 and we make them adjacent to v. For each such v where $\tau^{-1}(v) = \{z, w\}$, we set $\sigma(v_1) = z$ and $\sigma(v_2) = w$. In the resulting tree we pick an arbitrary edge, subdivide it and make the subdivision vertex adjacent to the new vertex r. In the resulting graph T', r will be the root. Let $f = (x, y)$ be an edge of T'. We set $X = \sigma(L(T_x)) \setminus \sigma(L(T_y))$, $Y = \sigma(L(T_y)) \setminus \sigma(L(T_x))$ and

$S = \sigma(L(T_x)) \cap \sigma(L(T_y))$. Notice that $\omega(f)$, as defined by the branch decomposition (T, μ), is exactly the set S, therefore S is a separator of G and $|S| \leq k$. We conclude that each edge in $\omega(f)$ has one endpoint in S and the other in Y. As each vertex in S has degree at most $\Delta(G)$, we obtain that $\omega(f) \leq k \cdot d$. Therefore, (T', σ) is an extended branch decomposition of G with width at most $k \cdot d$. The lemma now follows from Lemma 3. □

We are now ready to prove the main result of this paper.

Proof of Theorem 1. As $\mathcal{G}$ is non-trivial class we choose a graph $H \notin \mathcal{G}$ and we assume that H has h vertices. As $\mathcal{G}$ is minor closed, $\mathcal{G}$ excludes K_h as a minor. Therefore if $G \in \sqrt{\mathcal{G}}$, then G^2 does not contain K_h as a minor. From Lemma 1, G does not contain as a subgraph a properly subdivided wall of height $h - 2$. From Lemma 1, $\mathbf{bw}(G) < f(h-2)$. Notice also that $\Delta(G) \leq h-1$. Then, from Lemma 4, $\mathbf{cw}(G) < (h-1) \cdot f(h-2)$ and the theorem follows if we set $c_{\mathcal{G}} = (h-1) \cdot f(h-2)$. □

In the rest of this section, we will also bound the branch-width of the squares of the graphs in $\mathcal{G}$. This will be useful for the algorithmic consequences of our results in Section 4.

Lemma 5. *For every graph G it holds that* $\mathbf{cw}(G^2) \leq 2 \cdot (\mathbf{cw}(G))^2 - \mathbf{cw}(G)$.

Proof. Let (T, ρ) be a carving decomposition of G of width at most k. We will prove that (T, ρ) is a carving decomposition of G^2 of width at most $2k^2 - k$. We use the notation β and β' in order to distinguish the bridge sets for G and G^2 respectively. Let $f \in E(T)$. As $\Delta(G) \leq \mathbf{cw}(G)$, it is enough to prove that $\beta'(f) \leq 2 \cdot k \cdot (\Delta(G) - 1) + k$. To prove this, we first define $S = \rho^{-1}(T')$ where T' is one of the connected components of $T - \{e\}$. We define $L = S \cap (\bigcup_{e \in \beta(f)} e)$ and $R = \overline{S} \cap (\bigcup_{e \in \beta(f)} e)$. Finally, for each $e = \{x, y\} \in \beta(f)$ where $x \in L$ and $y \in R$, we define the sets

$$Q_R(e) = \{\{x, z\} \mid z \in N_G(y) \cap \overline{S}\}$$
$$Q_L(e) = \{\{z, y\} \mid z \in N_G(x) \cap S\}.$$

We claim that each edge $e' = \{v, u\} \in \beta'(f) \setminus \beta(f)$ belongs in $Q_R(e) \cup Q_L(e)$ for some $e \in \beta(f)$. For this, suppose that $e' = \{v, u\}$ where $v \in S$ and $u \in \overline{S}$. As e' is not an edge in G and there is a vertex $z \in V(G)$ such that $\{v, z\}, \{z, u\} \in E(G)$. If $z \in S$ then $\{z, u\} \in \beta(e)$ and therefore $e' \in Q_L(e)$ for $e = \{z, u\}$. If $z \in \overline{S}$ then $\{v, z\} \in \beta(e)$ and therefore $e' \in Q_R(e)$ for $e = \{v, z\}$. In both cases, the claim holds.

Notice now that for each $e = \{x, y\} \in \beta(f)$, $|Q_R(e) \cup Q_L(e)| = |Q_R(e)| + |Q_L(e)|$. As $|Q_R(e)| \leq |N_G(y)| - 1 \leq \Delta(G) - 1$ and $|Q_L(e)| \leq |N_G(x)| - 1 \leq \Delta(G) - 1$, we conclude that $|Q_R(e) \cup Q_L(e)| \leq 2 \cdot (\Delta(G) - 1)$. Therefore $|\beta'(f) \setminus \beta(f)| \leq 2 \cdot k \cdot (\Delta(G) - 1)$ and the claim follows. □

Lemma 6. *Let $\mathcal{G}$ be a class that is minor closed and let $G \in \sqrt{\mathcal{G}}$. Then* $\mathbf{bw}(G^2) \leq 4 \cdot c_{\mathcal{G}}^2 - 2 \cdot c_{\mathcal{G}}$.

Proof. Let G^2 be a graph in $\mathcal{G}$. From Theorem 1, $\mathbf{cw}(G) \leq c_{\mathcal{G}}$ and, from Lemma 5, $\mathbf{cw}(G^2) \leq 2 \cdot c_{\mathcal{G}}^2 - c_{\mathcal{G}}$. As G^2 is a graph, we have that $\Delta^*(G^2) = 2$. This together with Observation 2 and Lemma 4, implies that $\mathbf{bw}(G^2) \leq 4 \cdot c_{\mathcal{G}}^2 - 2 \cdot c_{\mathcal{G}}$. □

4. Linear time recognition of $\sqrt{\mathcal{G}}$

In this section, we study the complexity of the following problem when $\mathcal{G}$ is a minor closed graph class.

> Square Root of $\mathcal{G}$
> *Instance:* A graph G.
> *Question:* Does there exist a graph H such that H^2 is isomorphic to G?

Before we deal with the above problem, we briefly expose some known facts on the recognition of minor closed graph classes. Let $\mathcal{G}$ be such a class. We denote by $\mathbf{obs}(\mathcal{G})$ the set of minor minimal graphs that do not belong in $\mathcal{G}$. Notice that any two graphs in $\mathbf{obs}(\mathcal{G})$ are incomparable with respect to the minor relation. In their Graph Minors series, Robertson and Seymour [10] proved that each such set is finite, and therefore the size of its graphs is bounded by some constant $c'_{\mathcal{G}}$. Therefore, to check whether an input graph G is a member of some minor-closed $\mathcal{G}$, it suffices to check whether some graph on $\mathbf{obs}(\mathcal{G})$ is a minor of G. If this is the case, return NO as answer, otherwise return YES. In the same series of papers, Robertson and Seymour proved that to check whether a fixed size graph H is a minor of an n-vertex graph G can be done in $O(n^3)$ steps. Moreover, when the input graph G has branchwidth at most k, this check can be done in time $f(k) \cdot n$, that is linear in n (see e.g. [9,1,4]). We use this fact to prove our main algorithmic result.

Theorem 2. *For every non-trivial graph class $\mathcal{G}$ that is minor closed, there exists an algorithm deciding $\sqrt{\mathcal{G}}$ in linear time.*

Proof. Let $c_{\mathcal{G}}$ be the constant of Theorem 1 and let G be an input graph to the problem of the recognition of $\sqrt{\mathcal{G}}$. If $\Delta(G) > c_{\mathcal{G}}$, then, from Observation 1, $\mathbf{cw}(G) > c_{\mathcal{G}}$ and the answer to the problem is NO. Assuming that $\Delta(G) \leq c_{\mathcal{G}}$, the computation of G^2 can be done in linear time. From Lemma 6, if $G \in \sqrt{\mathcal{G}}$, then the branchwidth of G^2 is at most $c'_{\mathcal{G}} = 4 \cdot c_{\mathcal{G}}^2 - 2 \cdot c_{\mathcal{G}}$. We now check whether $\mathbf{bw}(G^2) \leq c'_{\mathcal{G}}$, e.g. using the algorithm of [2]. If the answer is negative, then we safely return NO. Otherwise, it remains to check whether $G \in \mathcal{G}$ for a graph of bounded branchwidth and this can be done in linear time as commented above. □

References

[1] Isolde Adler, Frederic Dorn, Fedor V. Fomin, Ignasi Sau, Dimtrios M. Thilikos, Faster parameterized algorithms for minor containment, in: 12th Scandinavian Workshop on Algorithm Theory—SWAT 2010 (Bergen), in: LNCS, vol. 6139, Springer, Berlin, 2010, pp. 322–333.
[2] Hans L. Bodlaender, Dimitrios M. Thilikos, Constructive linear time algorithms for branchwidth, in: 24th International Colloquium, Automata, Languages and Programming, ICALP'97, in: LNCS, vol. 1256, Springer, Berlin, 1997, pp. 627–637.
[3] Frank Harary, Richard M. Karp, William T. Tutte, A criterion for planarity of the square of a graph, J. Combin. Theory 2 (1967) 395–405.
[4] Illya V. Hicks, Branch decompositions and minor containment, Networks 43 (1) (2004) 1–9.
[5] Yaw Ling Lin, Steven S. Skiena, Algorithms for square roots of graphs, SIAM J. Discrete Math. 8 (1991) 99–118.
[6] László Lovász, Graph minor theory, Bull. Amer. Math. Soc. 43 (1) (2006) 75–86.
[7] Neil Robertson, Paul D. Seymour, Graph minors. V. Excluding a planar graph, J. Combin. Theory Ser. B 41 (1986) 92–114.
[8] Neil Robertson, Paul D. Seymour, Graph minors. X. Obstructions to tree-decomposition, J. Combin. Theory Ser. B 52 (2) (1991) 153–190.
[9] Neil Robertson, Paul D. Seymour, Graph minors. XIII. The disjoint paths problem, J. Combin. Theory Ser. B 63 (1) (1995) 65–110.
[10] Neil Robertson, Paul D. Seymour, Graph minors. XX. Wagner's conjecture, J. Combin. Theory Ser. B 92 (2) (2004) 325–357.
[11] Paul D. Seymour, Robin Thomas, Call routing and the ratcatcher, Combinatorica 14 (2) (1994) 217–241.