Josse van Dobben de Bruyn

Reduced divisors and gonality in finite graphs

Bachelor's thesis, 8 August 2012

Supervisor: Dr D.C. Gijswijt

Mathematisch Instituut, Universiteit Leiden

Contents

1 Introduction

In recent research, several results from the theory of Riemann surfaces have been related to similar results on graphs. For instance, the Riemann–Roch theorem for Riemann surfaces has an analogue for graphs, and the theory of (principal) divisors can also be applied to graphs. In that fashion, *gonality* has been defined for graphs as well. It turns out that the gonality of a graph is also related to winning strategies in certain chip-firing games on graphs, which have been studied since the 1980s.

In this thesis, the most important results and techniques are constructed in Section 3, most notably lower and upper bounds on the gonality. Some examples of graphs and their gonalities can be found in Section 4.

Finally we will conduct an in-depth analysis of an algorithm for computing the reduced divisor in Section 5. We also briefly discuss algorithms for other, related tasks, such as computing the rank of a divisor and computing the gonality of a graph.

This thesis gives an overview of some of the recent developments in the area of graph gonality, aimed at undergraduates in their final year as well as graduate students in mathematics. Some prior knowledge of graph theory and group theory might be useful, but in general the text is accessible for anyone with a general mathematical background (naive set theory, linear algebra).

2 Basic definitions

The theory of graph gonality is based upon various notions, which will briefly be discussed here.

2.1 Graph theory

Within this thesis, only finite, undirected, loopless multigraphs will be considered, which we will assume to be non-empty and connected. That is, a graph is a triple $G = (V, E, \phi)$ of finite disjoint sets V, E and a map $\phi : E \to [V]^2$, where $[V]^2$ denotes the set of 2-element subsets of V. The elements of the set V are called vertices and the elements of E are called edges. By definition, every edge $e \in E$ maps to a set $\{a, b\} \subset V$ consisting of exactly two distinct vertices, called the *ends* of e. For every $v \in V$, we let $E(v)$ denote the set of edges incident with v, that is: $E(v) = \{e \in E : v \in \phi(e)\}$. The degree $d(v)$ of a vertex $v \in V$ is the number of edges incident with v. In other words, we have $d(v) = |E(v)|$.

Throughout the majority of this text, we will only consider one graph at a time, which we will simply denote by G. Moreover we will call the vertex set V, the edges E and the edge map ϕ.

2.2 Divisors

A key concept in the theory of gonality is the concept of divisors.

Definition 2.1. The set $\mathrm{Div}(G)$ of *divisors on* G is the set of all functions $f : V \to \mathbb{Z}$. This becomes an abelian group by $(f + g)(v) := f(v) + g(v)$.

Definition 2.2. The *degree* $\deg(D)$ of a divisor $D \in \mathrm{Div}(G)$ is the sum of its entries:

$$\deg(D) = \sum_{v \in V} D(v).$$

For convenience, we make three more definitions.

Definition 2.3. For any $k \in \mathbb{Z}$, the set $\mathrm{Div}^k(G)$ is given by

$$\mathrm{Div}^k(G) = \{D \in \mathrm{Div}(G) : \deg(D) = k\}.$$

Note that $\mathrm{Div}^0(G)$ is a subgroup of $\mathrm{Div}(G)$. For any other $k \in \mathbb{Z}$, $\mathrm{Div}^k(G)$ is a coset of $\mathrm{Div}^0(G)$.

Definition 2.4. We define a partial order $\leq$ on $\mathrm{Div}(G)$ by writing $D \leq D'$ if and only if $D(v) \leq D'(v)$ holds for all $v \in V$. This uniquely determines the partial orders $<$, $\geq$ and $>$. Note that $D < D' \iff D \leq D' \wedge D \neq D'$, hence $D < D'$ doesn't imply that $D(v) < D'(v)$ holds for all $v \in V$; strict inequality only has to hold for some $v \in V$.

Definition 2.5. A divisor $D \in \mathrm{Div}(G)$ is said to be *effective* if $D \geq 0$. The set containing all effective divisors is denoted by $\mathrm{Div}_+(G)$. For any effective divisor $D \in \mathrm{Div}_+(G)$, we let $\mathrm{supp}(D)$ denote the *support* of D, that is:

$$\mathrm{supp}(D) = \{v \in V : D(v) > 0\}.$$

Definition 2.6. For all $k \in \mathbb{Z}$ we define

$$\mathrm{Div}_+^k(G) = \mathrm{Div}^k(G) \cap \mathrm{Div}_+(G).$$

That is, $\mathrm{Div}_+^k(G)$ is the set of all effective divisors of degree k on G.

2.3 The Laplacian matrix and principal divisors

One of the main elements of gonality theory is the Laplacian matrix of a graph, often simply referred to as the *Laplacian*. The easiest way to define the Laplacian of G is in terms of two other matrices A and D associated with G: D is the diagonal matrix given by $D_{v,v} = d(v)$ for all $v \in V$, and A is the adjacency matrix of G, which is given by

$$A_{v,w} = \begin{cases} \left|\{e \in E : \phi(e) = \{v, w\}\}\right| & \text{if } v \neq w, \\ 0 & \text{otherwise.} \end{cases}$$

It is obvious from the definition that A is symmetric.

Definition 2.7. The *Laplacian matrix* Q of G is given by $Q = D - A$. The homomorphism $\Delta : \mathrm{Div}(G) \to \mathrm{Div}(G)$ associated with the Laplacian matrix is called the *Laplace operator*.

Example 2.8. The complete bipartite graph $K_{2,3}$ has the following Laplacian matrix:

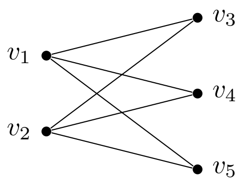

$$\begin{pmatrix} 3 & 0 & -1 & -1 & -1 \\ 0 & 3 & -1 & -1 & -1 \\ -1 & -1 & 2 & 0 & 0 \\ -1 & -1 & 0 & 2 & 0 \\ -1 & -1 & 0 & 0 & 2 \end{pmatrix}.$$

Note that the drawing of Q may depend on the order we choose to represent the vertices, however the abstract matrix Q is uniquely defined by the graph G. ⊠

Lemma 2.9. *Let $\mathbb{1}$ denote the all-ones vector. The Laplacian matrix Q, when viewed as a matrix over $\mathbb{Q}$, satisfies $\ker(Q) = \mathrm{span}(\mathbb{1})$ and $\mathrm{rank}(Q) = |V| - 1$*

Proof. As Q has zero row sums, we have $\mathbb{1} \in \ker(Q)$. For $|V| = 1$ the result is obvious, for now we have $\mathbb{Q}^V = \mathrm{span}(\mathbb{1}) \subset \ker(Q)$. Assume $|V| > 1$ and let $f \in \mathbb{Q}^V$ be given with $f \notin \mathrm{span}(\mathbb{1})$. We set $m = \max\{f(v) : v \in V\}$ and define $U \subset V$ by setting $U = \{v \in V : f(v) = m\}$. Note that $U \neq V$ holds because $f \notin \mathrm{span}(\mathbb{1})$. Because G is connected, we can choose some $u \in U$ which has a neighbour $w \in V \setminus U$. Now we have $A_{u,w} > 0$, so $A_{u,w} \cdot f(w) < A_{u,w} \cdot m$, hence

$$\begin{aligned} (Qf)(u) &= d(u) \cdot f(u) - \sum_{v \in V} (A_{u,v} \cdot f(v)) \\ &> d(u) \cdot f(u) - m \cdot \sum_{v \in V} A_{u,v} \\ &= d(u) \cdot m - m \cdot d(u) = 0. \end{aligned}$$

Hence $Qf \neq 0$. This holds for all $f \in \mathbb{Q}^V \setminus \mathrm{span}(\mathbb{1})$. Therefore we have $\ker(Q) = \mathrm{span}(\mathbb{1})$ and $\mathrm{rank}(Q) = |V| - 1$, which concludes the proof. □

Remark 2.10. Henceforth, we let Q denote the Laplacian matrix viewed as a matrix over any field of characteristic 0 (in particular $\mathbb{Q}$). On the other hand, Δ will exclusively be used to denote the homomorphism $\mathrm{Div}(G) \to \mathrm{Div}(G)$.

Definition 2.11. The divisors in the image of the homomorphism Δ are called *principal divisors*. The subgroup of principal divisors is denoted by $\mathrm{Prin}(G)$. In other words, we have

$$\mathrm{Prin}(G) = \Delta\left[\mathrm{Div}(G)\right].$$

Since every column of Q has total sum 0, every principal divisor has degree 0, so we have

$$\mathrm{Prin}(G) \subset \mathrm{Div}^0(G).$$

Since both $\mathrm{Prin}(G)$ and $\mathrm{Div}^0(G)$ are subgroups of $\mathrm{Div}(G)$, the above inclusion tells us that $\mathrm{Prin}(G)$ is in fact a subgroup of $\mathrm{Div}^0(G)$. This leads us to the next definition.

Definition 2.12. The *Jacobian* $\mathrm{Jac}(G)$ of G is defined to be the quotient group $\mathrm{Div}^0(G)/\,\mathrm{Prin}(G)$.

Theorem 2.13. *Let $\kappa(G)$ denote the number of spanning trees in G. The group $\mathrm{Jac}(G)$ is finite of order $\kappa(G)$.*

Proof. Choose some $v_0 \in V$. Now we define a map $\varphi : \mathbb{Z}^{V \setminus \{v_0\}} \to \mathrm{Div}^0(G)$ by

$$\varphi(D)(v) = \begin{cases} D(v) & \text{if } v \in V \setminus \{v_0\}, \\ -\sum\limits_{w \neq v_0} D(w) & \text{if } v = v_0. \end{cases}$$

That is, we set $\phi(D)(v_0)$ such that $\deg(\phi(D)) = 0$ holds. This is obviously an isomorphism. Let Q_{v_0} be the matrix obtained from Q by deleting the v_0-th row and column. Now the $|V|-1$ columns of Q_{v_0} span a lattice $\Lambda \subset \mathbb{Z}^{V \setminus \{v_0\}}$ of maximal rank. It follows that $\det(Q_{v_0}) = [\mathbb{Z}^{V \setminus \{v_0\}} : \Lambda]$. On the other hand, it follows from the matrix tree theorem that $\det(Q_{v_0}) = \kappa(G)$. Moreover, the deleted column is an integer multiple of the other columns (since G is connected), hence we have $\varphi[\Lambda] = \mathrm{Prin}(G)$. Therefore we have $[\mathrm{Div}^0(G) : \mathrm{Prin}(G)] = [\mathbb{Z}^{V \setminus \{v_0\}} : \Lambda] = \kappa(G)$. □

2.4 Linear equivalence and dimension

We define a relation $\sim$ on $\mathrm{Div}(G)$ by having $D \sim D'$ if and only if $D - D' \in \mathrm{Prin}(G)$. Note that this is an equivalence relation. This is often referred to as *linear equivalence* (e.g. in [2]), hence we will use that terminology as well.

Definition 2.14. The *linear system associated to* a divisor $D \in \mathrm{Div}(G)$ is defined to be the set $|D|$ containing all effective divisors that are equivalent to D, that is:

$$|D| = \{D' \in \mathrm{Div}_+(G) : D \sim D'\}.$$

The equivalence class of a divisor $D \in \mathrm{Div}(G)$ (with respect to linear equivalence) will be denoted by $[D]$. Note that $|D| = [D] \cap \mathrm{Div}_+(G)$.

Definition 2.15. The dimension $r(D)$ of the linear system $|D|$ is defined to be

$$r(D) = \max\left\{s \in \mathbb{Z} : |D - F| \neq \varnothing \text{ for all } F \in \mathrm{Div}_+^s(G)\right\}.$$

Remark 2.16. Note that the maximum is always attained, since $r(D) \leq \deg(D)$ holds. That is, for every $F \in \mathrm{Div}_+(G)$ with $\deg(F) > \deg(D)$ we have $\deg(D - F) < 0$, hence $|D - F| = \varnothing$. Moreover, observe that $r(D) = -1$ if and only if $|D| = \varnothing$.

Remark 2.17. Neither the linear system $|D|$ nor its dimension $r(D)$ depend on the choice of its representative D. That is, if $D \sim D'$, then obviously we have $|D| = |D'|$. Moreover, for all $F \in \mathrm{Div}(G)$ we have $D - F \sim D' - F$, hence $r(D) = r(D')$.

2.5 Gonality

As a special case of Remark 2.16, we have $\deg(D) \geq 1$ for all $D \in \mathrm{Div}(G)$ with $r(D) \geq 1$. Moreover, note that in every graph there is a divisor $D \in \mathrm{Div}(G)$ such that $r(D) \geq 1$; consider for instance the all-ones divisor $D = \mathbb{1}$. Now we can make the following definition.

Definition 2.18. The gonality $\mathrm{gon}(G)$ of a graph G is defined by

$$\mathrm{gon}(G) = \min \{\deg(D) : D \in \mathrm{Div}(G),\ r(D) \geq 1\}.$$

This minimum is guaranteed to exist, as the set $\{\deg(D) : D \in \mathrm{Div}(G),\ r(D) \geq 1\}$ is non-empty and bounded from below by 1.

Corollary 2.19. *For any graph G we have* $\mathrm{gon}(G) \geq 1$. □

Example 2.20. Let G be a tree. By Theorem 2.13 the Jacobian $\mathrm{Jac}(G)$ is trivial, so we have $\mathrm{Prin}(G) = \mathrm{Div}^0(G)$. Thus for any $D, D' \in \mathrm{Div}(G)$ with $\deg(D) = \deg(D')$ we have $D \sim D'$. In particular, for any $D, D' \in \mathrm{Div}_+^1(G)$ we have $0 \in |D - D'|$, so $r(D) \geq 1$. Now observe that $\mathrm{Div}_+^1(G)$ is not empty as for every $v \in V$ the unit divisor e_v, which is one in v and zero elsewhere, is an element of $\mathrm{Div}_+^1(G)$. Therefore we have $\mathrm{gon}(G) \leq 1$. By Corollary 2.19 we have $\mathrm{gon}(G) \geq 1$, hence $\mathrm{gon}(G) = 1$ must hold. ⊠

More examples will follow in Section 4. We will close this section with the following lemma, which is actually quite surprising. As it turns out, we can limit ourselves to effective divisors D in Definition 2.18. This is not immediately obvious from the definition, but it severely reduces the number of divisors we have to consider if we attempt to determine $\mathrm{gon}(G)$.

Lemma 2.21. *There exists some $D \in \mathrm{Div}_+(G)$ such that $\deg(D) = \mathrm{gon}(G)$ and $r(D) \geq 1$.*

Proof. Let $D' \in \mathrm{Div}(G)$ be a divisor such that $\deg(D') = \mathrm{gon}(G)$ and $r(D') \geq 1$. Now fix $v \in V$. Since $r(D') \geq 1$, there is some effective divisor $F \in \mathrm{Div}(G)$ such that $D' - e_v \sim F$. But now we have $D' - e_v - F \in \mathrm{Prin}(G)$, so it follows that $D' \sim F + e_v$. Since both F and e_v are effective, so is $F + e_v$. Thus we set $D = F + e_v$. Obviously, we have $D' \sim D$ and $\deg(D) = \deg(D') = \mathrm{gon}(G)$. Moreover, we can conclude from $D \sim D'$ that $r(D) = r(D') \geq 1$. □

Corollary 2.22. *We have* $\mathrm{gon}(G) = \min \{\deg(D) : D \in \mathrm{Div}_+(G),\ r(D) \geq 1\}$. □

3 Further techniques

3.1 Chip-firing games

The abstract concepts of the previous chapter can also be viewed in the context of chip-firing games. In a chip-firing game, we consider a graph G with any integer number of chips (commonly dollars, although other currencies are not at all uncommon) on its vertices, which can be moved about following certain rules. In the Norine–Baker chip-firing game ([2]), the rules are as follows:

1. The game is played by two people, you and an evil adversary. A graph G is given.
2. First, you have to divide a number of chips among the vertices of G. Vertices are allowed to have a negative number of chips; such vertices are said to be in debt.
3. The opponent then chooses any one vertex and subtracts one chip from that vertex. This vertex may already be in debt.
4. Subsequently, you are allowed to move the chips about. However, there is only one type of move you can do: you can move $\deg(v)$ chips from node v to all its neighbours by giving each of them 1 chip. You are not allowed to leave certain neighbours out or give some of them multiple chips! Using a sequence of moves of this type (called *firing moves* or *borrowing moves*), you have to make sure that all vertices are out of debt.
5. If you succeed, you win. Otherwise, the adversary wins.

We assume that the opponent plays his optimal strategy: he will win whenever he can. Thus, whether or not there will be a winning strategy for the remainder of the game depends on the choice of initial conguration in step 2.

Note that we can denote the state of G (that is, the number of chips on the vertices of G) at any moment during the game as a divisor on G. Now we have the following lemma, leading to an important theorem.

Lemma 3.1. *The firing moves correspond to subtracting principal divisors, and the initial divisor is equivalent to the resulting divisor. Moreover, every equivalent divisor can be reached via a finite sequence of firing moves.*

Proof. Let D be the initial divisor. Let e_v be the unit divisor corresponding to some vertex $v \in V$, then $-Qe_v$ is the vector containing $-\deg(v)$ on the v-th coordinate, $+1$ on each of the neighbours of v and 0 on all other vertices. Thus $D - Qe_v$ is the divisor resulting from firing from vertex v to all of its neighbours.

In order to see why every equivalent divisor can be reached using a finite sequence of firing moves, let D_1 be the initial divisor and D_2 an equivalent divisor. There exists some $f \in \mathrm{Div}(G)$ such that $D_2 = D_1 - \Delta f$, however the entries of f do not have to be positive. Now if we recall that $\mathbb{1} \in \ker(\Delta)$, we see that we can add any constant divisor $k \cdot \mathbb{1}$ to f in order to obtain $f' \in \mathrm{Div}_+(G)$ which also satisfies $D_2 = D_1 - \Delta f'$. Now f' corresponds to a finite sequence of firing moves, so every equivalent divisor can in fact be reached using a finite sequence of firing moves.[1] □

[1] The point is that the game only allows a positive number of firing moves on each vertex, whereas principal divisors may arise from a negative number of firing moves. The above argument assures us that even those principal divisors can be achieved within the rules of the chip-firing game.

Theorem 3.2. *The gonality of G equals the minimum number of chips in a winning initial configuarion in step 2 of the Norine–Baker chip-firing game.*

Proof. Assume that the game is winning for us if we start out with some divisor $D \in \mathrm{Div}(G)$. Subsequently, the opponent subtracts one chip from a vertex he chooses. Regardless of the chip he subtracts, there is a winning strategy for us (for the remainder of the game). That is: for every $F \in \mathrm{Div}_+^1(G)$, there is some effective divisor $F' \in \mathrm{Div}(G)$ such that $D - F \sim F'$. In other words: $|D - F| \neq \varnothing$. Hence we have $r(D) \geq 1$.

On the other hand, if the game is not winning for us after choosing some initial configuration $D \in \mathrm{Div}(G)$, then there must be some $F \in \mathrm{Div}_+^1(G)$ such that $|D - F| = \varnothing$, so now we have $r(D) < 1$. All in all, the game is winning after choosing some initial divisor $D \in \mathrm{Div}(G)$ if and only if $r(D) \geq 1$. □

As a result, we can think of gonality in terms of this chip-firing game. This helps our understanding of gonality and provides us with a useful toolkit for proving lemmata and theorems.

3.2 Firing from subsets

Apart from firing from a single vertex, one might also consider firing once from every vertex in some subset $A \subset V$. We will denote the indicator function of A as $\mathbb{1}_A$, that is:

$$\mathbb{1}_A(v) = \begin{cases} 1 & \text{if } v \in A, \\ 0 & \text{if } v \notin A. \end{cases}$$

Thus, henceforth the process of starting with some initial configuration $D \in \mathrm{Div}(G)$ and firing from a subset $A \subset V$ will be denoted by $D - \Delta\,\mathbb{1}_A$. When firing from a subset A, the edges between two vertices $v, w \in A$ will have one chip going from v to w and one chip going from w to v, so they cancel out. Hence we only have to worry about edges leaving A.

Definition 3.3. Let $A \subset V$ be a subset of vertices. The outdegree $\mathrm{outdeg}_A(v)$ of a vertex $v \in A$ with respect to A is the number of edges leaving A through v, that is:

$$\mathrm{outdeg}_A(v) = \Big|\big\{e \in E : \phi(e) \cap A = \{v\}\big\}\Big|.$$

Moreover, we let $\mathrm{outdeg}_A(A)$ denote the total outdegree of A. Similarly, the indegree $\mathrm{indeg}_A(v)$ of a vertex $v \in V \setminus A$ with respect to A is the number of edges entering $V \setminus A$ through v:

$$\mathrm{indeg}_A(v) = \Big|\big\{e \in E : \phi(e) \cap (V \setminus A) = \{v\}\big\}\Big|.$$

Remark 3.4. Note that the only difference is that outdeg only applies to vertices in A whereas indeg only applies to vertices in $V \setminus A$. That is, for all $v \in A$ we have $\mathrm{outdeg}_A(v) = \mathrm{indeg}_{V\setminus A}(v)$ and for all $v \in V \setminus A$ we have $\mathrm{indeg}_A(v) = \mathrm{outdeg}_{V\setminus A}(v)$.

Example 3.5. Assume that $D_2 = D_1 - \Delta\,\mathbb{1}_A$. For all $v \in A$ we have $D_2(v) = D_1(v) - \mathrm{outdeg}_A(v)$ and for all $v \in V \setminus A$ we have $D_2(v) = D_1(v) + \mathrm{indeg}_A(v)$. It is now apparent from $\Delta\mathbb{1} = 0$ that firing from a subset A can be undone by subsequently firing from $V \setminus A$. In particular, one could *reverse fire* into a vertex $v \in V$ by firing from every vertex in $V \setminus \{v\}$. Starting with an initial divisor $D \in \mathrm{Div}(G)$, we then end up with the divisor $D' = D + \Delta e_v$. This is also the main idea of the proof of Lemma 3.1. ⊠

Example 3.6. Let G be a tree and let $D \in \mathrm{Div}_+^1(G)$ be the effective divisor of degree 1 given by $D = e_v$ for some $v \in V$. For any neighbour w of v, the edge e with $\phi(e) = \{v, w\}$ induces a cut $(U, V \setminus U)$ separating v and w. By firing from U, we send one chip from v to w. All the firing moves on interior edges of U get leveled out by another firing move in the opposite direction, so the only thing that changes is the chip traveling from v to w. This corresponds with the finding that $\mathrm{Jac}(G)$ is trivial for any tree. ⊠

A major advantage of considering firing from subsets rather than firing from single vertices becomes apparent when we look at level sets.

Definition 3.7. Let $D, D' \in \mathrm{Div}(G)$ be two equivalent divisors with $D' = D - \Delta f$ for some $f \in \mathrm{Div}(G)$. Define $m = \max\{f(v) : v \in V\}$ and $k = m - \min\{f(v) : v \in V\}$. The *level set decomposition* of f is the sequence of sets $A_0 \subset A_1 \subset \cdots \subset A_k = V$ given by

$$A_i = \{v \in V : f(v) \geq m - i\}.$$

The sequence of divisors $D_0, D_1, \ldots, D_k \in [D]$ given by $D_0 = D$ and $D_{i+1} = D_i - \Delta\, \mathbb{1}_{A_i}$ for all $i \in \{0, 1, \ldots, k-1\}$ is called the *divisor sequence associated with the level set decomposition.*

Remark 3.8. Note that f is not uniquely determinded as we are free to add or subtract any integer multiple of $\mathbb{1}$, but the choice of f does not affect the level sets. Therefore we will usually assume that $f(v) \geq 0$ holds for all $v \in V$ in such a way that $f(v_0) = 0$ holds for at least one $v_0 \in V$. In that case we have $k = m$, so for all $i \in \{0, 1, \ldots, k\}$ we have $A_i = \{v \in V : f(v) \geq k - i\}$.

Lemma 3.9. *For all $i \in \{0, 1, \ldots, k-1\}$ we have $A_i \neq V$.*

Proof. By the definition of k, there is some $v \in V$ such that $f(v) = m - k$. Moreover, for all $i \in \{0, 1, \ldots, k-1\}$ and all $w \in A_i$ we must have $f(v) \geq m - i > m - k$, hence $v \notin A_i$. □

Theorem 3.10. *Let $D, D' \in \mathrm{Div}(G)$ be two equivalent divisors with $D' = D - \Delta f$ for some $f \in \mathrm{Div}(G)$. Let $A_0 \subset A_1 \subset \cdots \subset A_k$ be the level set decomposition of f and $D_0, D_1, \ldots, D_k$ the associated divisor sequence. Then for all $i \in \{0, 1, \ldots, k\}$ we have $D_i(v) \geq \min\big(D(v), D'(v)\big)$ (for all $v \in V$). Moreover, we have $D_k = D'$.*

Proof. Define $f_i \in \mathrm{Div}(G)$ recursively by $f_0 = f$ and $f_{i+1} = f_i - \mathbb{1}_{A_i}$. Inductively, it is not hard to see that $D' = D_i - \Delta f_i$ holds for all $i \in \{0, 1, \ldots, k\}$.

- For $i = 0$, we have $D_0 = D$ and $f_0 = f$, so we have $D' = D - \Delta f = D_0 - \Delta f_0$.
- Let $i \in \{0, 1, \ldots, k-1\}$ be given such that $D' = D_i - \Delta f_i$. Now note that $D_{i+1} = D_i - \Delta\, \mathbb{1}_{A_i}$ and $f_{i+1} = f_i - \mathbb{1}_{A_i}$, so we find

  $$D' = D_i - \Delta f_i = D_i - \Delta\left(f_{i+1} + \mathbb{1}_{A_i}\right) = D_i - \Delta\, \mathbb{1}_{A_i} - \Delta f_{i+1} = D_{i+1} - \Delta f_{i+1}.$$

  Hence the property holds for $j = i + 1$ as well.

Moreover, we will prove (once again by induction) that for all $i \in \{0, 1, \ldots, k\}$, any $v \in A_i$ satisfies $f_i(v) = m - i$, whereas any $v \in V \setminus A_i$ satisfies $f_i(v) < m - i$.

- For $i = 0$, this follows from the definition of A_i.
- Let $i \in \{0, 1, \ldots, k-1\}$ be given such that $f_i(v) = m - i$ holds for all $v \in A_i$ and $f_i(v) < m - i$ holds for all $v \in V \setminus A_i$. Let $v \in V$ be given. We distinguish three cases.

– If $v \in A_i$, then we have $f_i(v) = m - i$ by the induction hypothesis. Because we have $f_{i+1} = f_i - \mathbb{1}_{A_i}$, we must have $f_{i+1}(v) = m - i - 1$.

– If $v \in A_{i+1} \setminus A_i$, then we have $f_{i+1}(v) \leq f_i(v) < m - i$. Note that we have $v \notin A_j$ for all $j \leq i$, so we have $f_i(v) = f_0(v)$. Moreover, by definition we have $f_0(v) \geq m - i - 1$, so $f_{i+1}(v) \geq m - i - 1$. Combining these results, we find $f_{i+1}(v) = m - i - 1$.

– If $v \in V \setminus A_{i+1}$, then we have $f_0(v) < m - i - 1$. Since $f_{i+1} \leq f_0$, it follows that $f_{i+1}(v) < m - i - 1$.

All in all, we have $f_{i+1}(v) \leq m - (i + 1)$ for all $v \in V$, and equality holds if and only if $v \in A_{i+1}$. Thus the property holds for $j = i + 1$ as well.

Now we will prove that $D_i(v) \geq \min\big(D(v), D'(v)\big)$ holds for all $i \in \{0, 1, \ldots, k\}$ and all $v \in V$. Once again, we shall do so by induction on i.

- For $i = 0$, we have $D_0(v) = D(v) \geq \min\big(D(v), D'(v)\big)$ for all $v \in V$.

- Let $i \in \{0, 1, \ldots, k - 1\}$ be given such that $D_i(v) \geq \min\big(D(v), D'(v)\big)$ holds for all $v \in V$. Recall that $D_{i+1} = D_i - \Delta\,\mathbb{1}_{A_i}$. Now consider moving directly from D_i to D'. By the preceding discussion, we have $D' = D_i - \Delta f_i$. Moreover, for all $v \in A_i$ we have $f_i(v) = m - i$, whereas for all $v \in V \setminus A_i$ we have $f_i(v) < m - i$. Therefore, when moving directly from D_i to D', any $v \in A_i$ loses $(m - i) \cdot \text{outdeg}_{A_i}(v)$ chips whereas it receives at most $(m - i - 1) \cdot \text{outdeg}_{A_i}(v)$ chips. Thus for all $v \in A_i$ we have

$$D_i(v) - (m - i) \cdot \text{outdeg}_{A_i}(v) + (m - i - 1) \cdot \text{outdeg}_{A_i}(v) \geq D'(v),$$

  hence $D_i(v) \geq D'(v) + \text{outdeg}_{A_i}(v)$. Now that we know this, we fire from A_i in order to get to D_{i+1}. Let $v \in V$ be given. We distinguish two cases.

  – If $v \in A_i$, then we have $D_{i+1}(v) = D_i(v) - \text{outdeg}_{A_i}(v) \geq D'(v) \geq \min\big(D(v), D'(v)\big)$.

  – If $v \in V \setminus A_i$, then we have $D_{i+1}(v) \geq D_i(v) \geq \min\big(D(v), D'(v)\big)$.

  Hence for all $v \in V$ we have $D_{i+1}(v) \geq \min\big(D(v), D'(v)\big)$.

Finally, note that $A_k = V$ holds, so it follows from the preceding discussion that $f_k(v) = m - k$ holds for all $v \in V$. Therefore we have $f_k \in \ker(\Delta)$, so now it follows from $D' = D_k - \Delta f_k$ that $D' = D_k$. This concludes the proof. □

This means that we can go from D to D' via a sequence of subset-firing moves without ever going below the pointwise minimum of the two. One particularly useful consequence is the following.

Corollary 3.11. *Let $D, D' \in \text{Div}_+(G)$ be two equivalent (effective) divisors. Then there exists a finite sequence of subset-firing moves that transform D into D' without ever going into debt.*

Proof. This follows immediately from the previous theorem, as for all $v \in V$ we have $D(v) \geq 0$ and $D'(v) \geq 0$, hence $\min\big(D(v), D'(v)\big) \geq 0$. □

Remark 3.12. Note that no such sequence is guaranteed to exist if one merely considers firing from single vertices one at a time. For instance consider the following counterexample. Let $G = K_{1,k}$ for some $k \geq 2$ be a star centered around $c \in V$ and let $v, w \in V \setminus \{c\}$ be two vertices different from the center (see figure below). Then we set $D = e_v$ and $D' = e_w$. As we saw before in both Example 2.20 and Example 3.6, $\text{Jac}(G)$ is trivial, hence we have $D \sim D'$. Any intermediate divisor in a sequence of single chip-firing moves transforming D into D' without going into debt, contains only one chip. However at some point we have to pass through the center c in order to get from v to w. Unfortunately, firing from c would cause all of its neighbours (more than one) to receive a chip, hence c has to go into debt. Indeed, no sequence of single chip-firing moves exists transforming D into D' without ever going into debt.

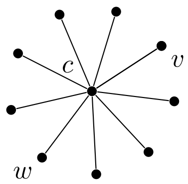

Example: the star $K_{1,9}$.

In Example 3.6 we already found a sequence of subset-firing moves transforming D into D' without ever going into debt, as the star $K_{1,k}$ is a tree. Corollary 3.11 assures us that such a sequence exists, indeed (however these two sequences are not guaranteed to be the same).

3.3 Reduction of a divisor

As Lemma 2.21 suggests, we can make quite some strong assumptions on our divisors. For instance, when looking for divisors $D \in \mathrm{Div}(G)$ of minimum degree with $r(D) \geq 1$, we can limit ourself to effective divisors. This severely reduces the number of divisors we have to consider, as $\mathrm{Div}_+^k(G)$ is finite for any $k \in \mathbb{Z}$ whereas $\mathrm{Div}^k(G)$ is infinite. A similar result is the reduction of divisors with respect to some vertex $v_0 \in V$.

Definition 3.13. For any $v_0 \in V$, a divisor $D \in \mathrm{Div}(G)$ is said to be v_0-*reduced* if both of the following conditions hold:

- $D(v) \geq 0$ for all $v \in V \setminus \{v_0\}$.
- For any non-empty $A \subset V \setminus \{v_0\}$ there is some $v \in A$ such that $D(v) < \mathrm{outdeg}_A(v)$.

Moreover, a divisor is said to be v_0-*semi-reduced* if it only satisfies the first property.

Reduced divisors can be interpreted as follows: every vertex other than v_0 is out of debt and we cannot fire from any non-empty subset (not containing v_0) without ruining the first property. The following results illustrate the importance of reduced divisors.

Lemma 3.14. *For any given $v_0 \in G$ and $D \in \mathrm{Div}(G)$, there exists a v_0-semi-reduced divisor $D' \in \mathrm{Div}(G)$ such that $D \sim D'$.*

Proof. We define $F \in \mathbb{Q}^V$ by setting $F = \mathbb{1} - |V| \cdot e_{v_0}$. Note that $F \in \mathrm{Div}^0(G)$ holds. As $\mathrm{Jac}(G)$ is finite, the order of $[F]$ in $\mathrm{Jac}(G)$ is finite as well. Hence there is some $k \in \mathbb{Z}_{>0}$ such that $[kF] = [0]$ holds. But now kF is a representative of $[0]$, hence a principal divisor. Thus we can add some multiple of kF to D in order to obtain a v_0-semi-reduced divisor $D' \sim D$: simply add kF as many times as needed to take all vertices other than v_0 out of debt. $\square$

Theorem 3.15 (Proposition 3.1 in [2]). *For any given $v_0 \in G$ and $D \in \mathrm{Div}(G)$, there exists a unique v_0-reduced divisor $D' \in \mathrm{Div}(G)$ such that $D \sim D'$.*

Proof. For any $v \in V$, let $l(v)$ denote the length of the shortest path between v and v_0. Moreover, we define $d \in \mathbb{Z}_{\geq 0}$ and $S_k \subset V$ for all $k \in \{0, 1, \ldots, d\}$ as follows:

$$d = \max_{v \in V} \ l(v);$$
$$S_k = \{v \in V : l(v) = k\}.$$

Now we define $\beta : \mathrm{Div}(G) \to \mathbb{Z}^{d+1}$ by setting

$$\beta(F) = \left(\sum_{v \in S_0} F(v), \sum_{v \in S_1} F(v), \ \ \ldots \ \ , \sum_{v \in S_d} F(v) \right).$$

Recall that there exists a v_0-semi-reduced divisor $D' \in [D]$ by Lemma 3.14. Now let $F \in [D]$ be a v_0-semi-reduced divisor such that

$$\beta(F) = \max \big\{ \beta(F') : F' \in [D] \text{ is } v_0\text{-semi-reduced} \big\}, \tag{1}$$

where we take the maximum with respect to the lexicographical order. Note that the maximum exists, as every coordinate is bounded by $\deg(D)$. (Moreover we have $\beta(F') \leq \big(\deg(D), 0, 0, \ldots, 0\big)$ for all $F' \in \mathrm{Div}(G)$ with $\deg(F') = \deg(D)$, so this holds in particular for every $F' \in [D]$ that are v_0-semi-reduced.)

Now assume that there is some non-empty subset $A \subset V \setminus \{v_0\}$ such that $F(v) \geq \mathrm{outdeg}_A(v)$ holds for all $v \in A$. We define $F' = F - \Delta(\mathbb{1}_A)$, that is: we fire from the entire subset A. Furthermore, we define $d_A = \min\{l(v) : v \in A\}$. For all $v \in A$ we have

$$F'(v) = F(v) - \mathrm{outdeg}_A(v) \geq 0, \tag{2}$$

so F' is v_0-semi-reduced as well. Now let $v' \in V$ be a node with $l(v') < d_A$ with a neighbour in A, for instance by taking the second node we encounter on a shortest path from A to v_0. (Such a node is guaranteed to exist, but it might be v_0). Now we have $F'(v') > F(v')$, since v' receives at least one chip as it has a neighbour in A. Moreover, for all $w \in V \setminus A$ with $l(w) < d_A$ we have $F'(w) \geq F(w)$, so we must have $\beta(F') > \beta(F)$. However this contradicts our choice of F. We can conclude that F already satisfies both of the properties from Definition 3.13, hence it is a v_0-reduced divisor. The choice $D' = F$ suffices for our main theorem.

In order to show that this v_0-reduced divisor is unique, suppose that both $D, D' \in \mathrm{Div}(G)$ are v_0-reduced with $D \neq D'$ and $D \sim D'$. Let $f \in \mathrm{Div}(G)$ be given with $D' = D - \Delta f$, let $A_0, A_1, \ldots, A_k$ be its level set decomposition and let $D_0, D_1, \ldots, D_k$ be the associated sequence of divisors. Note that $v_0 \in A_0$ must hold because D is v_0-reduced, so for all $i \in \{0, 1, \ldots, k\}$ we have $v_0 \in A_i$. Moreover, since $D \neq D'$ holds by assumption, we have $k \geq 1$. By Lemma 3.9 we have $A_{k-1} \neq V$, so we get from D_{k-1} to D' by firing from some subset $A_{k-1} \subsetneq V$ with $v_0 \in A_{k-1}$. Hence we can get from D' back to D_{k-1} by firing from $V \setminus A_{k-1}$, but this is a non-empty set not containing v_0. This contradicts the assumption that D' is v_0-reduced. Therefore we can conclude that there is only one v_0-reduced divisor $D' \sim D$. ☐

Because of the uniqueness of reduced divisors, we can make the following definition.

Definition 3.16. For any $D \in \mathrm{Div}(G)$ and $v_0 \in V$, we let $\mathrm{Red}_{v_0}(D) \in [D]$ denote the unique v_0-reduced divisor of D.

Lemma 3.17. *For any vertex $v_0 \in V$ and every v_0-semi-reduced divisor $D \in \text{Div}(G)$, we have $D'(v_0) \geq D(v_0)$, where $D' = \text{Red}_{v_0}(D)$ denotes the unique v_0-reduced divisor of D.*

Proof. This follows from our construction of D' in (1), as $S_0 = \{v_0\}$ holds. $\square$

The following theorem illustrates the importance of reduced divisors (and, to a lesser extent, the tools created in the proof of Theorem 3.15).

Theorem 3.18. *Let $D \in \text{Div}(G)$ and $v_0 \in G$ be given and let $D' = \text{Red}_{v_0}(D)$ be the v_0-reduced divisor of D. Then $r(D) \geq 0$ if and only if $D'(v_0) \geq 0$.*

Proof. If $r(D) \geq 0$, then there exists some $F \in \text{Div}_+(G)$ with $D \sim F$. Note that $F \sim D'$, so now it follows from Theorem 3.15 that $D' = \text{Red}_{v_0}(F)$. Moreover, note that F is already v_0-semi-reduced. It follows from Lemma 3.17 that $D'(v_0) \geq F(v_0) \geq 0$.

On the other hand, if $D'(v_0) \geq 0$, then D' is effective, so $D' \in |D|$ and $r(D) \geq 0$. $\square$

This theorem has a nice algorithmic application: finding the v_0-reduced divisor $\text{Red}_{v_0}(D)$ for some given $D \in \text{Div}(G)$ and $v_0 \in G$ yields an algorithm for determining whether or not $r(D) \geq 0$ holds. The steps in the proofs of Lemma 3.14 and Theorem 3.15 already provide an outline for the algorithm, but improvements in running time can be made. We will extensively study an algorithm for reducing divisors in Section 5.

Another nice little property of reduced divisors is given by the following lemma.

Lemma 3.19. *For any given $D \in \text{Div}(G)$, $v_0 \in V$ and $k \in \mathbb{Z}$, we have*

$$\text{Red}_{v_0}(D + k \cdot e_{v_0}) = \text{Red}_{v_0}(D) + k \cdot e_{v_0}.$$

Proof. Note that $\text{Red}_{v_0}(D)+k\cdot e_{v_0}$ is a v_0-reduced divisor, as $\text{Red}_{v_0}(D)$ is one. The result follows from the uniqueness of v_0-reduced divisors, since $\text{Red}_{v_0}(D) + k \cdot e_{v_0} \sim D + k \cdot e_{v_0}$ holds. $\square$

Theorem 3.20. *For any $D \in \text{Div}(G)$, we have $r(D) \geq 1$ if and only if $\text{Red}_v(D)(v) \geq 1$ holds for all $v \in V$.*

Proof. If $\text{Red}_v(D)(v) \geq 1$ holds for all $v \in V$, then we have $\text{Red}_v(D - e_v)(v) \geq 0$ for all $v \in V$ by Lemma 3.19, hence $r(D - e_v) \geq 0$ for all $v \in V$ by Theorem 3.18. Therefore we now have $|D - e_v| \neq \varnothing$ for all $v \in V$, so $r(D) \geq 1$.

On the other hand, if $r(D) \geq 1$ holds, then for every $v \in D$ we have $|D - e_v| \neq \varnothing$, hence $r(D - e_v) \geq 0$. Theorem 3.18 now assures us that $\text{Red}_v(D - e_v)(v) \geq 0$ holds for all $v \in V$. It follows from Lemma 3.19 that $\text{Red}_v(D)(v) \geq 1$ holds for all $v \in V$. $\square$

If we combine the theory of reduced divisors with level set decompositions, we find some more results. Firstly, there is a nice little lemma which assures us that we never have to fire from v_0 in order to reduce a v_0-semi-reduced divisor $D \in \text{Div}(G)$ with respect to v_0.

Lemma 3.21. *Let $D \in \text{Div}(G)$ be v_0-semi-reduced and let $D' = \text{Red}_{v_0}(D)$ be the reduced divisor of D. Furthermore, let $f \in \text{Div}_+(G)$ be such that $D' = D - \Delta f$ (with $f(w) = 0$ for some $w \in V$), and let $A_0, A_1, \ldots, A_k$ be the level set decomposition of f. Then for all $i \in \{0, 1, \ldots, k-1\}$ we have $v_0 \notin A_i$.*

Proof. Assume that there is some $i \in \{0, 1, \ldots, k-1\}$ such that $v_0 \in A_i$ holds. By definition we have $A_i \subset A_{k-1}$, hence $v_0 \in A_{k-1}$. Moreover, Lemma 3.9 assures us that $A_{k-1} \neq V$. Now let $D_0, D_1, \ldots, D_k$ be the sequence of divisors associated with the level set decomposition, satisfying $D_0 = D$ and $D_k = D'$. Now D_k is obtained from D_{k-1} by firing from some set $A_{k-1} \neq V$ which contains v_0. But then D_{k-1} can be obtained from D_k by firing from $V \setminus A_{k-1}$, which is a non-empty set not containing v_0. This contradicts the assumption that D' is v_0-reduced, hence we can conclude that $v_0 \notin A_i$ holds for all $i \in \{0, 1, \ldots, k-1\}$. $\square$

Finally, there is one more nice theorem with a topological interpretation, which uses level set decomposition as well as the above lemma. First we need to make one more definition.

Definition 3.22. For any vertex $v_0 \in V$ and any divisor $D \in \mathrm{Div}_+(G)$, the *v_0-component* $G_{v_0}(D)$ of D is the component containing v_0 in $G - \mathrm{supp}(D)$ if $v_0 \notin \mathrm{supp}(D)$, or the empty graph otherwise.

Example 3.23. In order to illustrate the above definitions, consider the following divisor.

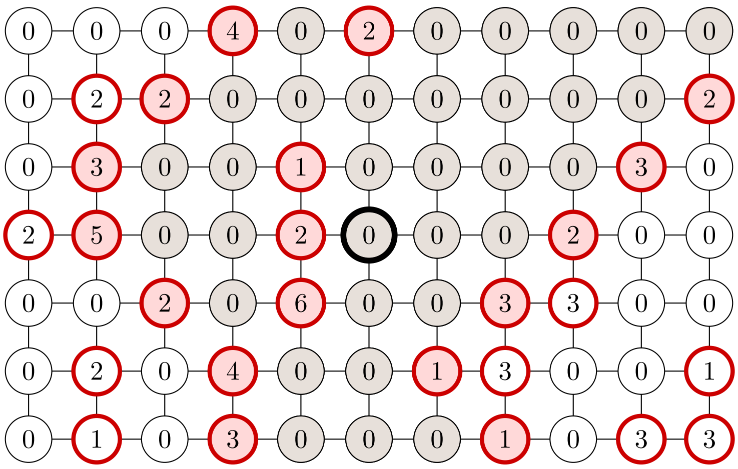

The v_0-component of D consists of the grey vertices and v_0 is highlighted by a thick black border. The vertices from $\mathrm{supp}(D)$ are highlighted by a thick red circle and those of them that bound $G_{v_0}(D)$ are also filled out in red. $\boxtimes$

Theorem 3.24. *Let $D, D' \in \mathrm{Div}_+(G)$ be given with $D' = \mathrm{Red}_{v_0}(D)$ and let $D_0, D_1, \ldots, D_k$ be the sequence of divisors associated with the level set decomposition (with $D_0 = D$ and $D_k = D'$). Then we have $G_{v_0}(D_j) \subset G_{v_0}(D_i)$ whenever $i \leq j$.*

Proof. It suffices to show that the inclusion holds whenever $j = i + 1$. If $G_{v_0}(D_{i+1})$ is empty, then the inclusion obviously holds. Moreover, if $G_{v_0}(D_i)$ is empty, then we have $D_i(v_0) = 0$, hence also $D_{i+1}(v_0) = 0$. It follows that $G_{v_0}(D_{i+1})$ is empty as well. For the remainder of this proof, we will assume that both $G_{v_0}(D_i)$ and $G_{v_0}(D_{i+1})$ are not empty.

Suppose, for the sake of contradiction, that there is some vertex $w_0 \in V$ with $w_0 \in V(G_{v_0}(D_{i+1}))$ and $w_0 \notin V(G_{v_0}(D_i))$. Let $P = w_0 w_1 w_2 \cdots w_l$ be some w_0-v_0-path in $G_{v_0}(D_{i+1})$ with $w_l = v_0$ (such a path is guaranteed to exist as $G_{v_0}(D_{i+1})$ is connected), and let w_r be the last node on P which does not belong to $G_{v_0}(D_i)$. Then we must have $D_i(w_r) > 0$ and $D_i(w_q) = 0$ for all $q \in \{r+1, r+2, \ldots, l\}$. Now we must have $w_r \in A_i$, because w_r loses at least one chip when we move from D_i to D_{i+1}. However w_{r+1} must also be in A_i, for otherwise it would receive at least one chip and give none away (which can not be true as w_{r+1} belongs to $G_{v_0}(D_{i+1})$). Inductively (by using the same exact argument as q increases) we see that $w_q \in A_i$ must hold for all $q \in \{r+1, r+2, \ldots, l\}$. In particular, we have $v_0 \in A_i$, but this contradicts the result of Lemma 3.21. It follows from this contradiction that every $w_0 \in G_{v_0}(D_{i+1})$ must also belong to $G_{v_0}(D_i)$, which concludes the proof. □

In terms of the reduction process, this means that $G_{v_0}(D_i)$ will only shrink as we move towards D_k using the respective firing moves given by level set decomposition. Moreover, the preceding theorem tells us that we can imagine reduction with respect to v_0 as a process of moving chips closer to v_0 in an attempt to make them land on v_0. (However this is not very accurate when e.g. $\mathrm{supp}(D) = V \setminus \{v_0\}$ holds, or when $D(v_0) > 0$ holds right from the start: in these cases the preceding theorem only assures us that $\varnothing \subset \varnothing$ holds.)

3.4 A lower bound on gon(G)

In most graphs, it is relatively easy to find an upper bound for the gonality: you have to find a divisor $D \in \mathrm{Div}(G)$ and prove that $r(D) \geq 1$. On the other hand, in order to find a lower bound $l \in \mathbb{Z}_{>0}$, you would have to prove that every divisor $D \in \mathrm{Div}(G)$ with $\deg(D) < l$ has $r(D) \leq 0$. For some specific graphs, that might be possible within a reasonable amount of time, however in most graphs this simply takes too long. Moreover, even when you find a lower bound, it is not guaranteed to be tight: you might just as well be way off. Hence it is desirable to develop a set of tools that provide us with lower bounds on the gonality.

Definition 3.25. A *bramble* is a non-empty set $\mathcal{B} \subset \mathbb{P}(V)$ with $\varnothing \notin \mathcal{B}$ and satisfying both of the following properties:

- For all $B \in \mathcal{B}$, the induced subgraph $G[B]$ is connected;
- For all $B, C \in \mathcal{B}$, we have $B \cap C \neq \varnothing$ or there is some edge $e \in E$ between B and C.

Thus, a bramble is a set of mutually touching connected vertex sets. The order $\|\mathcal{B}\|$ of $\mathcal{B}$ is the minimum cardinality of a cover of $\mathcal{B}$:

$$\|\mathcal{B}\| = \min \{|S| : S \subset V \text{ such that } S \cap B \neq \varnothing \text{ for all } B \in \mathcal{B}\}.$$

Definition 3.26. A *strict bramble* (or *strictly intersecting bramble*) is a bramble $\mathcal{B}$, which additionally satisfies the following, slightly stronger condition:

- For all $B, C \in \mathcal{B}$, we have $B \cap C \neq \varnothing$.

Definition 3.27. Let $D \in \mathrm{Div}_+(G)$ be a divisor and $\mathcal{B} \subset \mathbb{P}(V)$ a bramble. We define $\#_{\mathcal{B}}(D)$ and $\mathrm{m}_{\mathcal{B}}(D)$ by

$$\#_{\mathcal{B}}(D) = \big| \{B \in \mathcal{B} \mid \forall b \in B : D(b) = 0.\} \big|;$$

$$\mathrm{m}_{\mathcal{B}}(D) = \min \{\#_{\mathcal{B}}(D') : D' \in |D|\}.$$

That is, $\#_{\mathcal{B}}(D)$ is the number of connected vertex sets in the bramble that contain not a single chip at all. Note that the minimum $\mathrm{m}_{\mathcal{B}}(D)$ is guaranteed to exist as $\#_{\mathcal{B}}(D) \geq 0$ holds for every $D \in \mathrm{Div}_+(G)$.

Regarding strict brambles, we have the following theorem, which is based on unpublished work by M. Derickx.

Theorem 3.28. *Let $\mathcal{B}$ be a strict bramble. Then* $\text{gon}(G) \geq \|\mathcal{B}\|$.

Proof. Suppose that $\text{gon}(G) < \|\mathcal{B}\|$ holds. Then there exists some $D \in \text{Div}_+(G)$ with $r(D) \geq 1$ and $\deg(D) < \|\mathcal{B}\|$. Without loss of generality, we may assume that D is minimal in the sense that $\text{m}_{\mathcal{B}}(D) = \#_{\mathcal{B}}(D)$. Note that $\#_{\mathcal{B}}(D) \geq 1$ holds, as $\deg(D) < \|\mathcal{B}\|$. Hence there is some $B \in \mathcal{B}$ which contains no chips. Now we choose an arbitrary vertex $v_0 \in B$ and we let $D' = \text{Red}_{v_0}(D)$ be the v_0-reduced divisor of D. Corollary 3.11 (level set decomposition) gives us a finite sequence $D_0, D_1, \ldots, D_k$ of effective divisors which can be obtained by subsequently firing from subsets $A_i \subset V$ such that $D_0 = D$ and $D_k = D'$. Now let $i \in \{0, 1, \ldots, k\}$ be the smallest index such that B contains at least one chip in D_i, that is:

$$i = \min\big\{l \in \{0, 1, \ldots, k\} \mid \exists b \in B : D_l(b) > 0\big\}.$$

(Note that $D_k(v_0) > 0$ holds because D_k is v_0-reduced and $r(D) \geq 1$ holds.) By minimality of D, we have $\#_{\mathcal{B}}(D_i) \geq \#_{\mathcal{B}}(D)$. However, since B is now covered by at least one chip, there must be some other $C \in \mathcal{B}$ which is not covered anymore, but which was covered by the divisor D. Let D_j be the last divisor with $j < i$ in which C is still covered by a chip. Let A_j denote the set from which we fire to get from D_j to D_{j+1}. We will now prove the following two claims in order to get to a contradiction.

Claim: we must have $C \subset A_j$. If we were to have $w \notin A_j$ for some $w \in C$, then there must be some vertex $w' \in C$ with $w' \notin A_j$ which has a neighbour in A_j (as the induced subgraph $G[C]$ is connected). However, this vertex w' now receives at least one chip and does not lose a single chip, so we have $D_{j+1}(w') > 0$. This contradicts the assumption that C is not covered by any chip whatsoever in D_{j+1}. We can conclude that $C \subset A_j$ must hold.

Claim: we also have $B \subset A_j$. Note that $B \cap C \neq \varnothing$ holds, so we have $A_j \cap B \neq \varnothing$. Now assume that there is some $w \in B$ with $w \notin A_j$, then there must be some vertex $w' \in B$ with $w' \notin A_j$ which has a neighbour v' in $A_j \cap B$. Note that $D_j(v') = 0$ holds because $j < i$. However, now this vertex v' loses at least one chip when moving from D_j to D_{j+1}, because we have $\text{outdeg}_{A_j}(v') \geq 1$, but this contradicts the assumption that D_{j+1} is effective.

It follows from the previous discussion that $B \subset A_j$ must hold. However this contradicts Lemma 3.21, which assures us that we never have to fire from $v_0 \in B$. Therefore we can conclude that $\text{gon}(G) \geq \|\mathcal{B}\|$ must hold. □

Remark 3.29. Every graph contains a strict bramble, for instance $\mathcal{B} = \{V\}$. Unfortunately, this is not a very useful bramble, as we already knew that $\text{gon}(G) \geq 1$ holds.

Remark 3.30. In general, no strict bramble $\mathcal{B}$ with $\text{gon}(G) = \|\mathcal{B}\|$ is guaranteed to exist. For instance, consider the following counterexample: fix some $n \geq 4$ and let $G = K_n$ be the complete on n vertices. As we will see in Theorem 4.3, we have $\text{gon}(G) = n - 1$. Now suppose that there is some strict bramble $\mathcal{B}$ on G with $\|\mathcal{B}\| = n - 1$. Then for any $S \subset V$ with $|S| = n - 2$ there is a $B \in \mathcal{B}$ such that $S \cap B = \varnothing$ holds, but this means that $|B| \leq 2$. Now we choose $S' \subset V$ with $|S'| = n - 2$ such that $B \subset S'$ holds (here we use that $n \geq 4$, as otherwise such a set S' would not exist). Once again, there is some $B' \in \mathcal{B}$ such that $S' \cap B' = \varnothing$, but as $B \subset S'$ we now have $B \cap B' = \varnothing$. This is a contradiction because we assumed that $\mathcal{B}$ is a bramble. Therefore there does not exist a strict bramble $\mathcal{B}$ on G with $\text{gon}(G) = \|\mathcal{B}\|$ in this case.

3.5 An upper bound on gon(G)

We can also device a few tricks for finding an upper bound to the gonality.

Definition 3.31. A *strong separator* is a non-empty vertex set $S \subset V$ such that each component C in $G - S$ is a tree and every $s \in S$ satisfies $\big|\{e \in E(s) : \phi(e) \cap V(C) \neq \varnothing\}\big| \leq 1$.

In order to further illustrate this concept, consider the following example.

Example 3.32. Consider the Petersen graph P with the vertices of S drawn red.

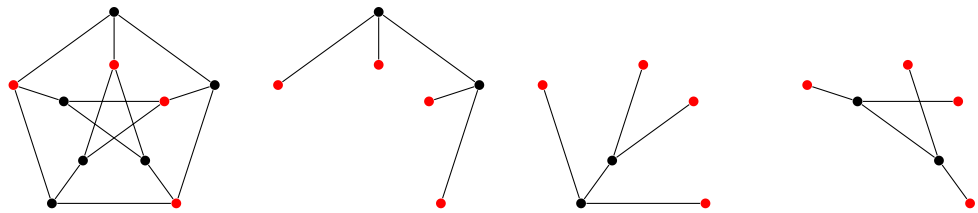

Now $P - S$ consists of three components, each of which is a path of length 1. Moreover, every $s \in S$ has exactly one edge going into each component, as can be seen in the figure. ⊠

Theorem 3.33. *Let $S \subset V$ be a strong separator, then* $\text{gon}(G) \leq |S|$.

Proof. Let $D \in \text{Div}(G)$ be given by $D = \mathbb{1}_S$. Choose any $v \in V$. If $v \in S$ holds, then $D - e_v$ is already effective, so now we obviously have $|D - e_v| \neq \varnothing$. Assume for the remainder of this proof that $v \notin S$ holds. Let C be the component of $G - S$ that v belongs to. Now let $D' = \text{Red}_v(D)$ be the v-reduced divsor of D and suppose that $D'(v) = 0$ holds. Then we let $H \subset G$ be the component containing v in the graph $G - \text{supp}(D')$, and let $B = V(H)$. By Theorem 3.24, H is an (induced) subgraph of the component in $G - \text{supp}(D)$ that v belongs to, which happens to be the component C. Now let $B' \subset \text{supp}(D')$ be given by

$$B' = \{b' \in \text{supp}(D') \mid \exists e \in E(b') : \phi(e) \cap B \neq \varnothing\}.$$

That is: B' consists of all the G-neighbours of B that lie outside B. Let $b' \in B'$ be given. We must now distinguish two cases in order to show that b' has *exactly* one neighbour in B.

- If $b' \in V(C)$, then we do not have to worry about multiple edges, because C is a tree. Suppose b' has two or more neighbours in B; let $u, w \in V(B)$ be two different neighbours of b'. Now as B is connected, both a v-u-path P_1 and a v-w-path P_2 must exist in B, but then there must be some cycle in C, which passes through b, u and w (given by $(P_1 \triangle P_2) \cup P_3$, where $\triangle$ denotes the edge-symmetric difference and P_3 is the path ubw). This contradicts the assumption that C is a tree, hence b' has only one neighbour in B.

- If $b' \notin V(C)$, then we must have $b' \in S$. Recall that $\big|\{e \in E(b') : \phi(e) \cap V(C) \neq \varnothing\}\big| \leq 1$, so b' has at most one incident edge going into C. Combine this with the fact that b' has at least one neighbour in $B \subset C$, so we can conclude that b' has exactly one neighbour in B.

All in all, we see that $\big|\{e \in E : \phi(e) \cap V(B) = \{b'\}\}\big| = 1$ holds for all $b' \in B'$. Therefore we have $\text{outdeg}_B(B) = |B'|$. But now we can fire from $V \setminus B$, without sending any vertex into debt, as only the vertices in B' lose chips (and each of them loses only one). However this contradicts the assumption that D' is v-reduced, as we can still fire from the non-empty set $V \setminus B$, which does not contain v, without going into debt. Hence $D'(v) \neq 0$ must hold. Recall from Lemma 3.17 that $D'(v) \geq D(v)$ holds, so now we must have $D'(v) > 0$. Therefore we have $|D - e_v| \neq \varnothing$, so $r(D) \geq 1$ and $\text{gon}(G) \leq |S|$. □

Remark 3.34. For any graph G the set $S = V$ is a strong separator. Moreover, for any simple vertex $v \in V$ (not incident with double edges) in a graph with two or more vertices, the set $S = V \setminus \{v\}$ is a strong separator.

Remark 3.35. Equality $\operatorname{gon}(G) = |S|$ does not have to hold, not even if we take a minimal strong separator in G. For instance, consider the following graph.

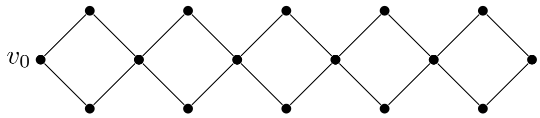

Now the divisor $D = 2e_{v_0}$ has rank (at least) one, as we can continue to move the chips on v_0 to the right by subsequently firing from all points from v_0 up to (and including) $\operatorname{supp}(D_i)$. The figure below illustrates this method. However, no two vertices form a strong separator, so we have $\operatorname{gon}(G) < |S|$ for every strong separator S.

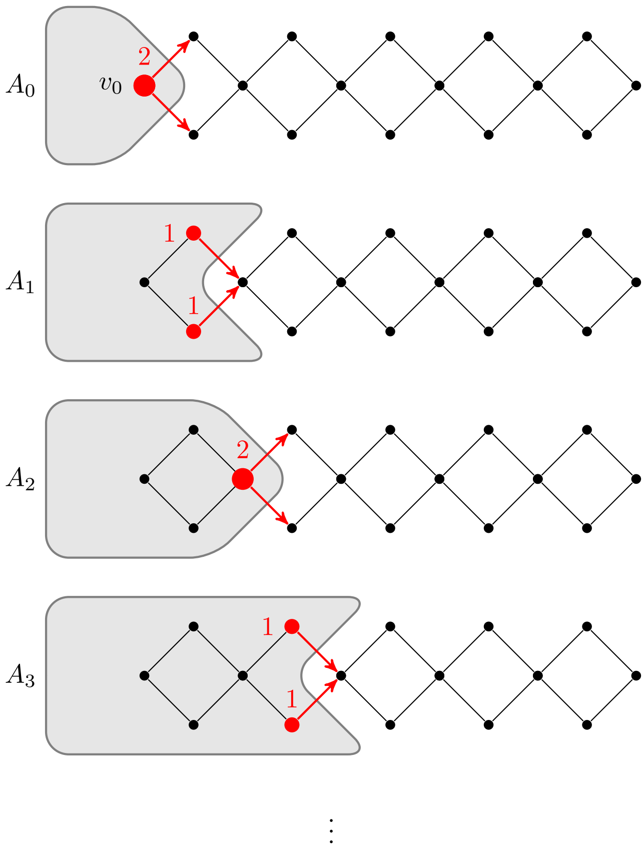

4 Examples

Now that we have constructed a set of tools, it is time to try and determine $\text{gon}(G)$ for some specific cases, including trees, cycles, graphs with only one cycle, complete graphs, complete bipartite graphs, the Petersen graph and two-dimensional grid graphs.

4.1 Trees

As we have seen on several occasions (Example 2.20 and Example 3.6), for any tree G the Jacobian is trivial and we have $\text{gon}(G) = 1$. As it turns out, trees are the only graphs with gonality zero.

Theorem 4.1. *Let G be a graph with $\text{gon}(G) = 1$. Then $\kappa(G) = 1$.*

Proof. There is some $D \in \text{Div}_+^1(G)$ with $r(D) \geq 1$. We can write $D = e_v$ for some $v \in V$. Since $r(D) \geq 1$, we have $e_v - e_w \in \text{Prin}(G)$ for all $w \in V$. Now let $F \in \text{Div}^0(G)$ be any divisor of degree zero, and set $F' = \text{Red}_v(F)$. Now we have $F'(w) \geq 0$ for all $w \in V \setminus \{v\}$ and $F'(v) = -\sum_{w \neq v} F'(w)$. It follows that

$$F = \sum_{w \in V \setminus \{v\}} F'(w) \cdot (e_w - e_v),$$

so we must have $F \in \text{Prin}(G)$. However this holds for all $F \in \text{Div}^0(G)$, so we must have $\text{Div}^0(G) = \text{Prin}(G)$. It follows from Theorem 2.13 that $\kappa(G) = 1$. $\square$

Corollary 4.2. *For any graph G we have $\text{gon}(G) = 1$ if and only if G is a tree.* $\square$

4.2 Cycles

The next graph we might want to consider, is the cycle C_n for $n \geq 2$. Because C_n is not a tree, it follows from the preceding corollary that $\text{gon}(C_n) \geq 2$. On the other hand, it is easy to see that any two vertices $v, w \in V$ with $v \neq w$ form a strong separator $S = \{v, w\}$, so it follows from Theorem 3.33 that $\text{gon}(C_n) \leq 2$. Therefore we have $\text{gon}(C_n) = 2$ for all $n \geq 2$.

Threorem 3.33 provides with a divisor $D \in \text{Div}_+(G)$ of rank at least one, given by $D = e_v + e_w$. Using the basic technique of the proof of said theorem, we see how we can move the two chips towards any chosen $v_0 \in V$. This is illustrated in the following figure.

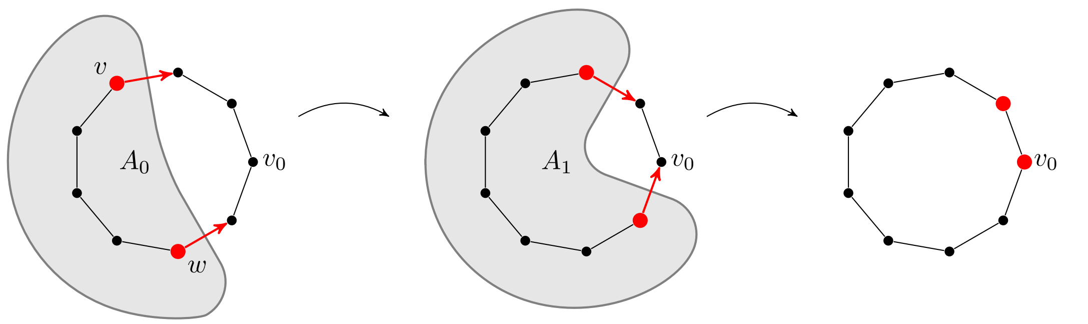

The same argument goes in fact to show that any graph containing only one cycle has gonality two: oviously any two vertices on that cycle form a strong separator, hence $\text{gon}(G) \leq 2$. Once again it follows from Corollary 4.2 that $\text{gon}(G) > 1$, hence we must have $\text{gon}(G) = 2$.

4.3 Complete graphs

The complete graph K_n allows a strong separator $S \subset V$ with $|S| = n-1$ for all $n \geq 2$ by Remark 3.34, hence for all $n \geq 2$ we have $\text{gon}(K_n) \leq n-1$. We can show that equality holds.

Theorem 4.3. *For all $n \geq 2$ we have $\text{gon}(K_n) = n-1$.*

Proof. For $n = 2$, equality obviously holds. If $n \geq 3$, we let $D \in \text{Div}_+(G)$ be any divisor of degree at most $n-2$. Now there are at least two vertices $v, w \in V$ such that $D(v) = D(w) = 0$ holds (but we will use only one of them). Note that D is v-semi-reduced. Let $A \subset V \setminus \{v\}$ be any non-empty subset not containing v and let $a \in \mathbb{Z}_{>0}$ be given by $a = |A|$. Now we have

$$\text{outdeg}_A(A) = a(n-a) = an - a^2.$$

Thus we have $\text{outdeg}_A(A) - n + 1 = an - a^2 - n + 1 = -(a-1)(a-n+1)$. Because $1 \leq a \leq n-1$ holds, we have $(a-1) \geq 0$ and $(a-n+1) \leq 0$, so we have $\text{outdeg}_A(A) - n + 1 \geq 0$. It follows that $\text{outdeg}_A(A) \geq n-1$. Now recall that the total number of chips in A is at most $n-2$, so we cannot fire from A without going into debt. As this holds for all non-empty $A \subset V \setminus \{v\}$, we see that D is v-reduced. It follows from Theorem 3.20 that $r(D) \leq 0$ must hold. This holds for every divisor of degree at most $n-2$, so we have $\text{gon}(K_n) \geq n-1$. Recall from the preceding remark that $\text{gon}(K_n) \leq n-1$, so in fact we have $\text{gon}(K_n) = n-1$. □

Remark 4.4. As a result, for every $k \in \mathbb{Z}_{>0}$ there is a graph of gonality k: simply consider the complete graph K_{k+1} on $k+1$ vertices.

4.4 Complete bipartite graphs

For any given $m, n \in \mathbb{Z}_{>0}$, consider the complete bipartite graph $K_{m,n}$. Let $V_1, V_2 \subset V$ denote the two parts of $K_{m,n}$. We will assume that $|V_1| \leq |V_2|$ holds. Now we define $D \in \text{Div}(G)$ by setting $D(v) = \mathbb{1}_{V_1}$. It is easy to see that $r(D) \geq 1$ holds, as for any $v \in V_2$ we have $D + \Delta e_v = \deg(D) \cdot e_v$. That is, we can simply fire from $V \setminus \{v\}$, as every $w \in V_1$ loses exactly one chip (which goes to v) and every $w \in V_2 \setminus \{v\}$ remains unchanged. Now $r(D) \geq 1$ holds, indeed. Therefore we have $\text{gon}(G) \leq \min(m, n)$. Once again, equality holds.

Theorem 4.5. *For any given $m, n \in \mathbb{Z}_{>0}$ we have $\text{gon}(K_{m,n}) = \min(m, n)$.*

Proof. This goes analogously to the proof of Theorem 4.3. If $\min(m, n) = 1$ holds, then equality obviously holds, so assume that $\min(m, n) > 1$. We will prove that $\text{outdeg}_A(A) \geq \min(m, n)$ holds for all $A \subset V$ with $0 < |A| < |V|$. Without loss of generality, we can assume that $m \leq n$ holds, hence $|V_1| = m$ and $|V_2| = n$. We define $a = |V_1 \cap A|$ and $b = |V_2 \cap A|$. Now we have

$$\text{outdeg}_A(A) = a(n-b) + b(m-a).$$

Obviously we have $0 \leq a \leq m$ and $0 \leq b \leq n$. If either $a = 0$ or $n - b = 0$ holds (or both), then we must have $b > 0$ and $m - a > 0$. In both cases this yields $\text{outdeg}_A(A) \geq m$. Similarly, if $b = 0$ or $m - a = 0$ holds (or both), then we must have $a > 0$ and $n - b > 0$, so now we have $\text{outdeg}_A(A) \geq m$ as well. Finally, if the four of them (a, b, $n-a$ and $n-b$) are all nonzero, then we have $\text{outdeg}_A(A) \geq a + (m-a) = m$. Thus, we have $\text{outdeg}_A(A) \geq m$ in any case.

Let $D \in \text{Div}_+(G)$ be any divisor of degree at most $\min(m, n) - 1$. There is some $v \in V$ with $D(v) = 0$, because $\deg(D) < m + n$. However we cannot fire from any non-empty $A \subset V \setminus \{v\}$ without going into debt, as $\text{outdeg}_A(A) \geq m$ holds for all $A \subset V$ with $0 < |A| < |V|$. Hence D is v-reduced, so it follows from Theorem 3.20 that $r(D) \leq 0$. This holds for all $D \in \text{Div}_+(G)$ with $\deg(D) < \min(m, n)$, so we have $\text{gon}(G) \geq \min(m, n)$. □

4.5 Two-dimensional grids

Let $G_{m,n}$ for $m, n \in \mathbb{Z}_{>0}$ be the two-dimensional $m \times n$ grid graph. We can assume that $m \leq n$ holds, for otherwise we could simply interchange m with n. We can define a divisor $D \in \mathrm{Div}_+(G)$ satisfying $\deg(D) = m$ by setting $D(v) = 1$ for every $v \in V$ in the leftmost column of the grid and $D(v) = 0$ for every other v. We have $r(D) \geq 1$ because we can move a chip to every vertex in the following way:

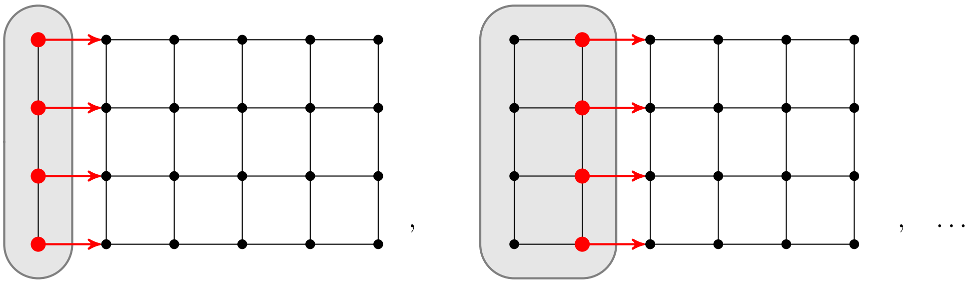

Therefore, we have $\mathrm{gon}(G_{m,n}) \leq m$. Now let $R_i \subset V$ for $i \in \{1, 2, \dots, m\}$ denote the i-th row and let C_j for $j \in \{1, 2, \dots, n\}$ denote the j-th column. We define the sets $B_{i,j} \subset V$ by setting $B_{i,j} = R_i \cup C_j$. For example, the subset $B_{3,5} \subset G_{4,6}$ is given in the figure below.

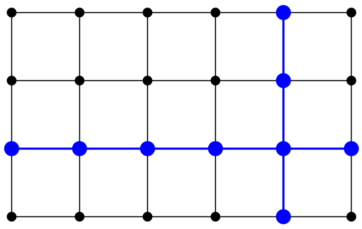

All the $B_{i,j}$ are connected. Moreover, for any $i, k \in \{1, 2, \dots, m\}$ and any $j, l \in \{1, 2, \dots, n\}$ we have $\{v_{i,l}, v_{k,j}\} \subset B_{i,j} \cap B_{k,l}$, so in particular the intersection is not empty. Therefore the set $\mathcal{B} = \{B_{i,j} : i \in \{1, 2 \dots, m\},\ j \in \{1, 2, \dots, n\}\}$ is a strict bramble.

Suppose that $\|\mathcal{B}\| < m$. Then there is some $S \subset V$ such that $S \cap B_{i,j} \neq \varnothing$ holds for all $i \in \{1, 2, \dots, m\}$ and all $j \in \{1, 2, \dots, n\}$. Since $|S| < m \leq n$ holds, there must be some row R_i such that $S \cap R_i = \varnothing$ and some column C_j such that $S \cap C_j = \varnothing$. But now the set $B_{i,j} = R_i \cup C_j$ satisfies $B_{i,j} \cap S = \varnothing$: we have a contradiction. Therefore we must have $\|\mathcal{B}\| \geq m$. Indeed, when we take $S = C_1$, then we clearly cover every $B_{i,j}$. In both cases, this set S satisfies $|S| = m$, so we have $\|\mathcal{B}\| = m$. Now it follows from Theorem 3.28 that $\mathrm{gon}(G_{m,n}) \geq m$.

All in all, we have seen that $\mathrm{gon}(G_{m,n}) \leq m$ and $\mathrm{gon}(G_{m,n}) \geq m$, so we must have equality. That is: $\mathrm{gon}(G_{m,n}) = m$ holds for all $m, n \in \mathbb{Z}_{>0}$ with $m \leq n$. Thus for any $m, n \in \mathbb{Z}_{>0}$ (not necessarily satisfying $m \leq n$) we have $\mathrm{gon}(G_{m,n}) = \min(m, n)$.

The divisor D can be generalised to higher dimensional grids, so we have $\mathrm{gon}(G) \leq |G| \,/ \max_i m_i$ for every d-dimensional grid graph G of size $m_1 \cdot m_2 \cdots m_d$. However a strict bramble of such high order is not guaranteed to exist. (For instance, if we take $B_{i,j,k}$ analogously in a three dimensional grid, then not every pair has a common vertex.) Thus we have the following conjecture.

Conjecture 4.6. *For every grid graph G of size $m_1 \cdot m_2 \cdots m_d$, we have* $\mathrm{gon}(G) = \dfrac{|G|}{\max\limits_{i \leq d} m_i}$.

5 Algorithmic approach

In Section 4 we saw that the gonality can be established for certain classes of graphs that are sufficiently nice in one way or another, but in general it is quite a difficult task to establish the gonality of any given graph with the tools developed so far. Therefore it might be worth our while to look for an algorithm to determine the gonality of a given graph. Unfortunately, no polynomial time algorithm is known to exist. However we can make life a little easier by providing algorithms for related problems.

5.1 Divisor reduction

The reduction of a divisor $D \in \mathrm{Div}(G)$ with respect to some base vertex $v_0 \in V$ has already proven itself to be a useful representative of D in the equivalence class $[D]$. In order to find the v_0-reduced divisor $D' \in [D]$, consider the following pseudo-algorithm.

1. Find a v_0-semi-reduced divisor $D' \sim D$ such that $0 < D'(v) < 2d(v)$ for all $v \in V \setminus \{v_0\}$.
2. Find some non-empty set $A \subset V \setminus \{v_0\}$ with $D(v) \geq \mathrm{outdeg}_A(v)$ for all $v \in A$ and fire from that particular subset (replace D by $D - \Delta\,\mathbb{1}_A$). Repeat this step until no such A exists.

We will elaborate on both steps in detail.

Algorithm 5.1 (semi-reduction)**.** The first step of the reduction algorithm is executed as follows.

1. Define $x \in \mathrm{Div}(G)$ by setting $x(v) = d(v) - D(v)$ for all $v \in V \setminus \{v_0\}$ and choosing $x(v_0)$ such that $\deg(x) = 0$ holds. (Recall that $d(v)$ denotes the degree of v.)
2. Now find some $y \in \mathbb{Q}^V$ with $y(v_0) = 0$ such that $x = Qy$ holds. Note that there is a unique $y \in \mathbb{Q}^V$ which satisfies this property, as we have $\ker(Q) = \mathrm{span}\,\{\mathbb{1}\}$. This can be done by Gaussian elimination (which can be done in polynomial time, see [8]).
3. Finally, we define $D' \in [D]$ by setting $D' = D + \Delta\,\lfloor y \rfloor$, where $\lfloor y \rfloor$ is obtained by rounding down all the entries in y.

The point of this first step in the main reduction algorithm is not only to find a semi-reduced divisor $D' \in [D]$, but also to limit the values of D' in order to reduce the number of firing moves we need to make in the second part of the reduction algorithm. The following lemma quantifies the bounds on D'.

Lemma 5.2. *Let $D' \in [D]$ be obtained by Algorithm 5.1. Then for all $v \in V \setminus \{v_0\}$ we have*

$$0 < D'(v) < 2d(v). \tag{3}$$

Proof. Define $z \in \mathbb{Q}^V$ by setting $z = y - \lfloor y \rfloor$ and note that $0 \leq z(v) < 1$ holds for all $v \in V$. Moreover, we have

$$D' + Qz = D + \Delta\,\lfloor y \rfloor + Q\,(y - \lfloor y \rfloor) = D + Qy = D + x,$$

hence by our definition of x we have $(D' + Qz)\,(v) = d(v)$ for all $v \in V \setminus \{v_0\}$. Moreover, for all $v \in V \setminus \{v_0\}$ we have

$$(Qz)(v) = -d(v)z(v) + \sum_{\substack{e \in E \\ \phi(e) = \{v,w\}}} z(w).$$

Because $0 \le z(u) < 1$ holds for all $u \in V$ (and $d(v) > 0$), it follows that

$$(Qz)(v) \ge -d(v)z(v) > -d(v).$$

Moreover, we use $0 \le z(u) < 1$ to find

$$(Qz)(v) < -d(v)z(v) + \sum_{\substack{e \in E \\ v \in \phi(e)}} 1 = -d(v)z(v) + d(v) \le d(v).$$

Thus for all $v \in V \setminus \{v_0\}$ we have $D'(v) = d(v) - (Qz)(v) > 0$ and $D'(v) = d(v) - (Qz)(v) < 2d(v)$, which concludes the proof. □

An outline for the remainder of the algorithm has already been provided in the proof of Theorem 3.15: find a set $A \subset V \setminus \{v_0\}$ such that $D(v) \ge \text{outdeg}_A(v)$ holds for all $v \in A$ (and fire from that set) until no such set exists. However it is not immediately obvious how we can proceed to find such a set in polynomial time, as there are $2^{|V|-1}$ subsets of $V \setminus \{v_0\}$ to consider. However, as we will see, the following algorithm will do.

Algorithm 5.3 (further reduction)**.** We split the second step into two pieces.

(a) Find a non-empty subset $A \subset V \setminus \{v_0\}$ such that we can fire from A without going into debt. That is, we have to assure that $D(v) \ge \text{outdeg}_A(v)$ holds for all $v \in A$. We will do this by building a sequence $A_0, A_1, \ldots, A_k$ of sets obtained by subsequently removing a single vertex. In doing so, we can easily keep track of the indeg and outdeg of every vertex in $V \setminus A_i$ and A_i, respectively. We color the vertices: every vertex is white as long as it is included in A_i, it becomes grey when it is scheduled to be removed and it becomes black as soon as it has actually been removed from A_i.

PART 2A: FINDPOSSIBLESET($G = (V, E, \phi)$, $D \in \text{Div}(G)$, $v_0 \in V$)

Require: $D(v) \ge 0$ for all $v \in V \setminus \{v_0\}$

```
A_0 ← V
i ← 0
ENQUEUE(q, v_0)
for all v ∈ V do
    state[v] ← WHITE
    indeg[v] ← 0
    outdeg[v] ← 0
end for
while not QUEUE-EMPTY(q) do
    v ← DEQUEUE(q)
    A_{i+1} ← A_i \ {v}
    i ← i + 1
    state[v] ← BLACK
    for all neighbours w of v do
        j ← |{e ∈ E : φ(e) = {v, w}}|
        if w ∈ A_i then
            indeg[v] ← indeg[v] + j
            outdeg[w] ← outdeg[w] + j
            if state[w] = WHITE and outdeg[w] > D[w] then
                ENQUEUE(q, w)
                state[w] ← GREY
```

22. **end if**
23. **else**
24. indeg$[w] \leftarrow$ indeg$[w] - j$
25. **end if**
26. **end for**
27. **end while**
28. $k \leftarrow i$
29. **return** $(A_k, \text{indeg}, \text{outdeg})$

(b) If the final set A_k is empty, the algorithm terminates. Otherwise, we fire from A as many times as we can (that is, $\min\{D(v)/\operatorname{outdeg}_A(v) : v \in A,\ \operatorname{outdeg}_A(v) > 0\}$ times) and start over with part 2a. The following procedure puts it all together.

PART 2B: FURTHERREDUCE($G = (V, E, \phi)$, $D \in \mathrm{Div}(G)$, $v_0 \in V$)

Require: $D(v) \geq 0$ for all $v \in V \setminus \{v_0\}$

repeat
  $(A, \text{indeg}, \text{outdeg}) \leftarrow$ FINDPOSSIBLESET(G, D, v_0)
  if $A \neq \varnothing$ **then**
    $l \leftarrow \min\{D[v]/\operatorname{outdeg}[v] : v \in A$ **and** $\operatorname{outdeg}[v] > 0\}$
    for all $v \in V$ **do**
      if $v \in A$ **then**
        $D[v] \leftarrow D[v] - l \cdot \operatorname{outdeg}[v]$
      else
        $D[v] \leftarrow D[v] + l \cdot \operatorname{indeg}[v]$
      end if
    end for
  end if
until $A = \varnothing$
return D

Lemma 5.4. *If* FINDPOSSIBLESET *returns a set $A_k \subset V$, then we have $v_0 \notin A_k$. Moreover, the variable* outdeg$[v]$ *corresponds to the actual value of* $\operatorname{outdeg}_{A_k}(v)$ *for all $v \in A_k$.*

Proof. Firstly, observe that v_0 is enqueued in line 3, so it is guaranteed to be dequeued and removed in lines 10–11. In order to see that $\operatorname{outdeg}[v] = \operatorname{outdeg}_{A_i}[v]$ holds at the end of each iteration of the while loop (lines 9–27), we use induction on i.

- For $i = 0$, we have $A_0 = V$, hence $\operatorname{outdeg}_{A_0}(v) = 0$ for all $v \in A_0$. This is indeed correct.
- Let $i \in \{0, 1, \ldots, k\}$ be given such that the statement is true for all $j \in \{0, 1, \ldots, i\}$. In iteration $j + 1$, we set $A_{j+1} = A_j \setminus \{v\}$, where v comes fresh off the queue. Since we only removed one vertex from A_j, for all $w \in A_{j+1}$ we have

  $$\operatorname{outdeg}_{A_{j+1}}(w) = \operatorname{outdeg}_{A_j}(w) + \big|\{e \in E : \phi(e) = \{v, w\}\}\big|.$$

  For every neighbour $w \in A$ of v, this corresponds to the amount we are adding in line 18. Moreover, for every non-neighbour $w \in A$, we have $\big|\{e \in E : \phi(e) = \{v, w\}\}\big| = 0$, so outdeg$[w]$ should remain unchanged. Indeed it does.

In particular, we have $\operatorname{outdeg}_{A_k}(v) = \operatorname{outdeg}[v]$ by the end of FINDPOSSIBLESET. □

Remark 5.5. Analogously, the variable indeg$[v]$ corresponds to the actual $\operatorname{indeg}_{A_k}(v)$ for all $v \in V \setminus A_k$ by the end of the procedure. In order to make this work, we have to increase indeg[v] once for every neighbour that is still in A_k (this happens when v is removed from A_i) and decrease indeg[w] if we find a neighbour w which has already been excluded from A_i.

Lemma 5.6. *If* FindPossibleSet *returns a non-empty set A_k, then for all $v \in A_k$ we have $D(v) \geq \operatorname{outdeg}_{A_k}(v)$.*

Proof. Assume that there is some $w \in A_k$ such that $\operatorname{outdeg}_{A_k}(v) > D(v)$ holds by the end of FindPossibleSet. Note that w cannot have been enqueued, as we remove elements that have been enqueued from the set A_i at some point. Now in particular there is some point in time where we increment outdeg[w] to become larger than $D(w)$, because initially we had outdeg[w] $= 0 \leq D(w)$. The only increment is made in line 18, so this must have happened at that time. Immediately afterwards, outdeg[w] $> D[w]$ holds, so either state[w] $=$ White and we enter into lines 20–21 or we must have had state[w] $\neq$ White. Note that w has never been enqueued, so apparently we must have had state[w] $\neq$ White in line 19. However line 21 is the only line where elements that have not been dequeued yet change state, but this happens immediately after Enqueue(q, w). This is a contadiction, as w has not been enqueued, so we can conclude that $D(v) \geq \operatorname{outdeg}_{A_k}(v)$ holds for all $v \in A_k$. ☐

Lemma 5.7. *Let $A_k \subset V \setminus \{v_0\}$ be a set returned by* FindPossibleSet*. If there exists some non-empty subset $B \subset V \setminus \{v_0\}$ such that $D(b) \geq \operatorname{outdeg}_B(b)$ holds for all $b \in B$, then $B \subset A_k$.*

Proof. Suppose that $B \not\subset A_k$ holds. Let $i \leq k$ be the smallest index such that $B \not\subset A_{i+1}$ and let $b \in B$ be the element that is to be removed from A_i. Note that $\operatorname{outdeg}_{A_j}(v) \leq \operatorname{outdeg}_B(v)$ holds for all $j \leq i$, hence in particular we have $D(b) \geq \operatorname{outdeg}_{A_j}(b)$ for all $j \leq i$. It follows from (the proof of) Lemma 5.4 that outdeg[b] corresponds with the actual value of $\operatorname{outdeg}_{A_j}(b)$ at every time $j \leq i$ during the algorithm, so at any stage of the algorithm before b is removed we have $D[b] \geq$ outdeg[b]. But now b has never passed the test in line 19, hence has never been enqueued in line 20. This is a contradiction, so we can conclude that $B \subset A_k$ must hold. ☐

Remark 5.8. It follows from Lemma 5.7 (combined with Lemma 5.6) that the set A_k returned by FindPossibleSet is the maximal set $M \subset V \setminus \{v_0\}$ satisfying $D(v) \geq \operatorname{outdeg}_M(v)$ for all $v \in M$. Moreover, for all $v \in V \setminus (A_k \cup \{v_0\})$ we must have $D(v) < d(v)$, for otherwise we would have $\{v\} \subset A_k$ by Lemma 5.7.

Corollary 5.9. *The procedure* FindPossibleSet *will return a non-empty set $A_k \subset V \setminus \{v_0\}$ such that $D(v) \geq \operatorname{outdeg}_{A_k}(v)$ holds for all $v \in A_k$ if and only if such a set exists.*

Proof. If such a set exists, then FindPossibleSet will return a non-empty set by Lemma 5.7. It follows from Lemma 5.6 that A_k satisfies the required property (recall that $v_0 \notin A_k$ holds). On the other hand, if no such set exists, then A_k must be empty by Lemma 5.6. ☐

This assures us that the procedure FindPossibleSet does exactly what it is supposed to do. We will look at the correctness of FurtherReduce later on, but first we will establish some bounds on the running time of FindPossibleSet.

Lemma 5.10. *The lines 10–13 in* FindPossibleSet *will be executed at most $|V|$ times and the lines 15–25 will be executed at most $2 \cdot |E|$ times.*

Proof. Note that every vertex is enqueued at most once, hence the lines 10–13 will be executed at most once for each vertex. Moreover, the lines 15–25 will be executed $d(v)$ times for every v that is dequeued in line 10, so these lines are executed at most

$$\sum_{v \in V} d(v) = 2|E|$$

times. If we use edge lists with multiplicities to represent the graph G (that is, we have a list of neighbours for each vertex as well as a list of multiplicities of the edges), then these lines are executed at most $2 \cdot \|E\|$ times, where $\|E\| = \left|\phi[E]\right|$ is the number of pairs $\{v, w\} \in [V]^2$ that are connected by one or more edges. $\square$

In order to determine the total running time of the algorithm, we will henceforth assume that arithmetic operations can be executed in constant time.

Corollary 5.11. *Using an efficient implementation, the procedure* FindPossibleSet *can run in $\mathcal{O}(\|E\|)$ worst-case running time.*

Proof. All the operations in lines 10–26 can be implemented to run in constant time. For instance, we might save A_i as an array of bits indicating whether or not the vertex at index v is an element of A_i. Queueing and dequeueing can be done in constant time as well, and the value of j can simply be a constant in an edge list-implementation of multigraphs, where each edge is bundled with a cardinality (rather than storing seperate copies for different edges sharing the same ends). All the other operations can obviously be executed in constant time. Note that $\|E\| \geq |V| - 1$, as G is connected, hence we have $\mathcal{O}(\|E\| + |V|) = \mathcal{O}(\|E\|)$ worst-case running time. $\square$

Now we have a look at the procedure FurtherReduction, which puts it all together.

Theorem 5.12. *The procedure* FurtherReduction *will find the unique v_0-reduced divisor $\mathrm{Red}_{v_0}(D)$ given any v_0-semi-reduced divisor D and base vertex $v_0 \in V$.*

Proof. We saw in the proof of Theorem 3.15 that firing from a vertex set $A \subset V \setminus \{v_0\}$ yields an equivalent divisor $D' \in [D]$ with $\beta(D') > \beta(D)$. It follows from the proof of Theorem 3.15 that the number of steps we have to take in this fashion is finite, hence FurtherReduction terminates. It follows from Corollary 5.9 that such a set exists if and only if the set A_k returned by FindPossibleSet is not empty. Thus, the divisor D is v_0-reduced if and only if A is empty in line 3 (and subsequently in line 13) of FurtherReduction. $\square$

Now that we have seen that FindPossibleSet runs in $\mathcal{O}(\|E\|)$ time, it would be nice to find some upper bound for the number of times we have to call that procedure in order to finish the reduction process. Indeed, when we first use Algorithm 5.1 in order to find a v_0-semi-reduced divisor $D' \sim D$ (which satisfies the property from Lemma 5.2), then we can devise a bound on the number of times we have to call the procedure FindPossibleSet in line 2 of FurtherReduction.

Lemma 5.13. *Let D' be the divisor producer by Algorithm 5.1, let $F = \mathrm{Red}_{v_0}(D')$ and let $f \in \mathrm{Div}_+(G)$ be such that $F = D' - \Delta f$ and $f(v_0) = 0$. Then $\deg(f) \leq 4 \cdot |E| \cdot |V| \cdot (|V| - 1)$.*

Proof. Let $A_0, A_1, \ldots, A_k$ be the level set decomposition of f and let $B_0, B_1, \ldots, B_l$ be the subsequence of $A_0, A_1, \ldots, A_{k-1}$ consisting of the sets that are actually different from one another, in their respective order (note that we exclude A_k because $A_k = V$). Now we have

$$f = \mathbb{1}_{A_0} + \mathbb{1}_{A_1} + \ldots + \mathbb{1}_{A_{k-1}} = c_0\, \mathbb{1}_{B_0} + c_1\, \mathbb{1}_{B_1} + \ldots + c_l\, \mathbb{1}_{B_l},$$

where each c_i denotes the number of times that B_i occurs among $A_0, \ldots, A_{k-1}$. Now when we subsequently fire c_i times from B_i, then every intermediate divisor F_i is an element of the sequence $D_0, D_1, \ldots, D_k$ of divisors associated with the level set decomposition. Therefore, in particular, every intermediate divisor is v_0-semi-reduced. As every B_i loses at least one chip every time it fires, it follows that

$$\begin{aligned} c_i &\leq \sum_{v \in B_i} F_i(v) \leq \sum_{v \in V \setminus \{v_0\}} F_i(v) = \deg(F_i) - F_i(v_0) \\ &\leq \deg(D) - D(v_0) = \sum_{v \in V \setminus \{v_0\}} D'(v) \leq \sum_{v \in V \setminus \{v_0\}} 2 \cdot d(v) \leq 4 \cdot |E|. \end{aligned}$$

As we have $\varnothing \subsetneq B_0 \subsetneq B_1 \subsetneq \cdots \subsetneq B_l \subsetneq V$, we must have $l \leq |V| - 2$. Therefore we have

$$\max\{f(v) : v \in V\} \leq \sum_{i=0}^{l} c_i \leq 4 \cdot |E| \cdot (l+1) \leq 4 \cdot |E| \cdot (|V| - 1).$$

As $f \geq 0$, it follows that $\deg(f) \leq 4 \cdot |E| \cdot |V| \cdot (|V| - 1)$. □

Corollary 5.14. *Let D' be obtained by Algorithm 5.1. Then the procedure* FINDPOSSIBLESET *is executed at most $4 \cdot |E| \cdot |V| \cdot (|V| - 1)$ times within the procedure* FURTHERREDUCE*(G, D', v_0).* □

Corollary 5.15. *Let D' be obtained by Algorithm 5.1. The worst-case running time of the procedure* FURTHERREDUCE*(G, D', v_0) is $\mathcal{O}(|V|^2 |E| \cdot \|E\|)$.* □

Therefore, we can reduce any divisor in polynomial time. Note that the matrix inversion from Algorithm 5.1 has a lower worst-case time complexity than the final reduction from Algorithm 5.3, so the total worst-case running time for reducing a divisor is $\mathcal{O}(|V|^2 |E| \cdot \|E\|)$.

Another upper bound on the number firing moves in FURTHERREDUCE is found by F. Shokrieh in [9], who devised an algorithm very similar to the algorithm presented here.

Theorem 5.16 (of F. Shokrieh)**.** *Let λ_2 denote the smallest eigenvalue of the matrix Q_{v_0}, which is obtained from Q by removing both the v_0-th row and column. Then the total number of vertex firing moves in* FURTHERREDUCTION *is at most $\frac{8|E|\sqrt{|V|}}{\lambda_2}$.*

5.2 Rank of a divisor

Theorem 3.18 and Theorem 3.20 immediately provide us with algorithm for deciding whether $r(D) \geq 0$ and $r(D) \geq 1$ hold, respectively.

Algorithm 5.17. In order to check whether $r(D) \geq 0$ and $r(D) \geq 1$ hold, consider the following algorithms.

- DIMENSIONNONNEGATIVE$(G = (V, E), D \in \mathrm{Div}(G))$
  Choose some $v_0 \in V$.
  $D \leftarrow$ SEMIREDUCE(G, D, v_0)
  $D \leftarrow$ FURTHERREDUCE(G, D, v_0)
  if $D(v_0) \geq 0$ **then**
      return true
  else
      return false
  end if

- DIMENSIONPOSITIVE($G = (V, E), D \in \mathrm{Div}(G)$)

  for all $v_0 \in V$ **do**

  $F \leftarrow$ SEMIREDUCE(G, D, v_0)

  $F \leftarrow$ FURTHERREDUCE(G, F, v_0)

  if $F(v_0) \leq 0$ **then**

  return false

  end if

  end for

  return true

Remark 5.18. Corectness of these algorithms is assured by Theorem 3.18, Theorem 3.20 and Theorem 5.12. Moreover, by Corollary 5.15, the worst-case running time of the former is $\mathcal{O}(|V|^2|E| \cdot \|E\|)$, whereas the worst-case running time of the latter is $\mathcal{O}(|V|^3|E| \cdot \|E\|)$.

5.3 Gonality

In order to compute the gonality of a given graph, it suffices to run DIMENSIONPOSITIVE(G, D) for every $D \in \mathrm{Div}_+(G)$ with $\deg(D) \leq |V|$, or even just $\deg(D) \leq |V| - 1$ if G contains a simple vertex (this follows from Remark 3.34), and then return the lowest degree divisor such that DIMENSIONPOSITIVE returns **true**. In order to get a grasp of the amount of divisors we would have to check, we determine the number of elements in $\mathrm{Div}_+^k(G)$ for all $k \in \mathbb{Z}$.

Lemma 5.19. *For any $k \in \mathbb{Z}$ we have* $\left|\mathrm{Div}_+^k(G)\right| = \binom{|V| + k - 1}{|V| - 1}$.

Proof. When $k = 0$, then $\mathrm{Div}_+^k(G) = \{0\}$, hence the statement holds. For $k < 0$ the statement is trivial, as $\mathrm{Div}_+^k(G) = \varnothing$ and $\binom{|V|+k-1}{|V|-1} = 0$. Thus we assume that $k > 0$ holds. Write every divisor $D \in \mathrm{Div}_+^k(G)$ as a string of $|V| + k$ letters v and c, where v represents a new vertex and c represents a chip on the preceding vertex. The first letter always has to be v, so we could just as well leave it out. Now we have an injective map $\psi : \mathrm{Div}_+^k(G) \to \{v, c\}^{|V|+k-1}$ such that $\psi(D)$ contains exactly $|V| - 1$ times letter v and k times letter c for any $D \in \mathrm{Div}_+^k(G)$. On the other hand, every string which contains $|V| - 1$ times letter v and k times letter c results from some divisor D as we can easily reconstruct D from this string. Therefore $\left|\mathrm{Div}_+^k(G)\right|$ equals the number of said sequences, which yields the result. $\square$

Remark 5.20. This is also known in combinatorics as the problem of putting k indistinguishable balls into n distinguishable boxes.

Thus, computing the gonality in the manner just explained can not always be done in polynomial time, that is: the algorithm has superpolynomial worst-case complexity.

Conjecture 5.21. *There does not exist a polynomial time algorithm for calculating* $\mathrm{gon}(G)$.

6 Further reading

Baker and Norine provide a nice proof of the Riemann–Roch theorem for graphs in [2], as well as some other interesting results. The Riemann–Roch theorem for graphs uses the cannonical divisor $K \in \mathrm{Div}(G)$ given by $K(v) = d(v) - 2$ and states the following.

Theorem 6.1. *Let G be a graph, and let D be a divisor on G. Then*

$$r(D) - r(K - D) = \deg(D) + 1 - g,$$

where $g = |E| - |V| + 1$ is the genus of G.

This is in fact the graph theoretic analogue of the Riemann–Roch theorem for Riemann surfaces. Moreover, they define the Abel–Jacobi maps $S^k_{v_0} : \mathrm{Div}(G) \to \mathrm{Jac}(G)$ by

$$S^k_{v_0}(D) = \Big[D - k \cdot e_{v_0}\Big].$$

This leads to the following two theorems.

Theorem 6.2. *The map S^k is surjective if and only if $k \geq g$. (As in the previous theorem, g denotes the genus of G.)*

Theorem 6.3. *The map S^k is injective if and only if G is $(k+1)$-edge-connected.*

Hladký, Král and Norine ([5]) relate divisors on a graph G to its corresponding divisors on the corresponding metric graph Γ, in which every edge has length one.

Theorem 6.4. *Let D be a divisor on a graph G and let Γ be the metric graph corresponding to G. Then, $r_G(D) = r_\Gamma(D)$.*

This leads them to the following corollary, which is all discrete.

Theorem 6.5. *Let D be a divisor on a graph G and let G^k be the graph obtained from G by subdividing each edge of G exactly k times. The ranks of D in G and in G^k are the same.*

Their paper also established a connection between divisors on tropical curves and metric graphs. Other works in this direction include [1] by Baker, [6] by Luo and [7] by van der Pol. The results are interesting, but cannot easily be expressed in the context of discrete (that is, regular, non-metric) graphs.

Björner, Lovász and Shor ([3]) study a chip-firing game first introduced by J. Spencer in 1986 ([10]). In terms of divisors, this game can be described as follows.

- We start with an initial divisor $D \in \mathrm{Div}_+(G)$ on a finite, simple, loopless graph, which is assumed to be non-empty and connected.
- The only moves allowed are firing from a single chip such that it does not go into debt.

Among their results are the following.

Theorem 6.6. *Given a connected graph and an initial distribution of chips, either every legal game can be continued infinitely, or every legal game terminates after the same number of moves with the same final position. Moreover, the number of times a given node is fired is the same in every legal game.*

Theorem 6.7. *Let $N = \deg(D)$ denote the number of chips, $n = |V|$ the number of vertices and $m = |E|$ the number of edges.*

(a) If $N > 2m - n$ then the game is infinite.

(b) If $m \leq N \leq 2m - n$ then there exists an initial configuration guaranteeing finite termination and also one guaranteeing infinite game.

(c) If $N < m$ then the game is finite.

For further reading on graph theory, [4] is a very nice textbook.

6.1 A conjecture

In every graph we encountered so far, the gonality equals the tree width. A well known property of brambles (which is proven in [4]) is the following.

Theorem 6.8 (Seymour & Thomas 1993)**.** *Let $k \geq 0$ be an integer. A graph has tree width $\geq k$ if and only if it contains a bramble of order $> k$.*

This leads us to the following conjecture.

Conjecture 6.9. *If $\mathcal{B}$ is a bramble, then* $\operatorname{gon}(G) \geq \|\mathcal{B}\| - 1$.

This would imply that $\operatorname{gon}(G) \geq \operatorname{tw}(G)$. As we have seen in Theorem 3.28, a very similar result holds if we only allow strictly intersecting brambles. However the proof of said theorem does not extend to brambles, as for instance any tree on $|V| \geq 2$ vertices allows a bramble of order 2.

References

[1] Matthew Baker: *Specialization of linear systems from curves to graphs*, Algebra & Number Theory, 2 (2008) 613–653.

[2] Matthew Baker and Serguei Norine: *Riemann–Roch and Abel–Jacobi theory on a finite graph*, Adv. Math. 215 (2007), no. 2, 766–788.

[3] A. Björner, L. Lovász, P.W. Shor, *Chip-firing games on graphs*, European Journal of Combinatorics 12 (4) (1991) 283–291.

[4] Reinhard Diestel: *Graph Theory* (fourth edition), Springer-Verlag, 2010.

[5] Jan Hladký, Daniel Král and Serguei Norine: *Rank of divisors on tropical curves*, `http://arxiv.org/abs/0709.4485v3`.

[6] Ye Luo, *Rank-determining sets of metric graphs*, `http://arxiv.org/abs/0906.2807`.

[7] Jorn van der Pol, *Analysis of the Brill-Noether game on metric cactus graphs*, `http://alexandria.tue.nl/extra1/afstversl/wsk-i/pol2011.pdf`.

[8] A. Schrijver, *Theory of linear and integer programming*, Wiley, 1986.

[9] Farbod Shokrieh: *Chip-firing games, G-parking functions, and an efficient bijective proof of the matrix-tree theorem*, `http://arxiv.org/abs/0907.4761`.

[10] J. Spencer: *Balancing vectors in the max norm*, Combinatorica 6 (1986), 55–66.