Distribution Agreement

Griffin Lee Miller April 1, 2019

On Direct-Sum Decompositions of the Picard Group of a Graph

by

Griffin Lee Miller

David Zureick-Brown
Adviser

Department of Mathematics

David Zureick-Brown

Adviser

Dwight Duffus

Committee Member

Susan Bredlau

Committee Member

2019

On Direct-Sum Decompositions of the Picard Group of a Graph

By

Griffin Lee Miller

David Zureick-Brown

Adviser

An abstract of
a thesis submitted to the Faculty of Emory College of Arts and Sciences
of Emory University in partial fulfillment
of the requirements of the degree of
Bachelor of Sciences with Honors

Department of Mathematics

2019

Abstract

On Direct-Sum Decompositions of the Picard Group of a Graph
By Griffin Lee Miller

The Picard group is an algebraic object that naturally arises in the study of a "chip-firing game," a solitaire game played on the vertices of a graph. This Picard group has illuminated the study of multiple research areas, offering a link between disparate topics in combinatorics, graph theory, and algebra. Recently, algebraic geometers have taken interest in the game, as it pertains to the developing subfield "tropical geometry." Such an exchange allows for the import of chip-firing observations into a geometric setting and has produced graph-theoretic analogues of classical algebraic geometry results. This thesis investigates the structure of the Picard group, particularly its behavior under a "coning" operation that we can iterate on its referent graph. Coning over a graph G entails adding a vertex to its vertex set which is adjacent to every other vertex in G. Recent papers demonstrate that coning over a graph produces a class of chip-firing game configurations that correspond to a subgroup of this Picard group. We investigate when the Picard group of the nth cone over G is the direct sum of this subgroup and another subgroup defined by the remaining game configurations. In particular, we show conditions on n for this decomposition to hold for a handful of different graphs, we show that this decomposition holds for infinitely many n on any graph, and we indicate a graph-theoretic property which necessitates this decomposition holds. This last result suggests there are graph properties that can elucidate the algebraic structure of the Picard group, which calls for further investigation.

On Direct-Sum Decompositions of the Picard Group of a Graph

By

Griffin Lee Miller

David Zureick-Brown

Adviser

A thesis submitted to the Faculty of Emory College of Arts and Sciences
of Emory University in partial fulfillment
of the requirements of the degree of
Bachelor of Sciences with Honors

Department of Mathematics

2019

Acknowledgements

I would first like to thank my advisor, David Zureick-Brown, for providing me with a deep, entertaining topic to study for the past year and for the hours of advisement provided in the process of researching and writing this thesis. Dr. Zureick-Brown has been a profound and challenging mentor who has guided me through a course of study that I look back on fondly.

I would also like to thank Dr. Duffus for providing comments on this work and for sitting on my committee. Thanks to Dr. Bredlau for sitting on my committee as well and for being a professor who left an indelible impact on my education. Thank you also to Dr. Parimala, whose course in Algebra was highly enlightening.

I am grateful for my father, the perspective he offered at crucial points in my life, and our conversations about academia, music, and New York. I am grateful for my mother, who invigorated my passion for writing and has acted as a role model and confidant, and my step-father, who has been a vital pillar in my development.

Thank you, Emma, for being a wonderful sister and fellow Sagittarian. You are the biggest hater and I hope that your ears close.

Thank you, Elise, for being my partner in all of this. It goes without saying that I would not be the person I am without your presence in my life. I will never express my gratitude enough! Crazy to see the journey we've taken over these past four years and to remember the teenagers who decided they wanted to come down to Atlanta. You the realest!!

I want to close by thanking all of the friends who have stuck with me through this crazy life! Shout out everyone who held it down in New York while I was gone. Shout out Bill, Vasi, Katie, Julia, Farris, Robby, Mustafa, Clara, Luis, Simon, and the many others who have sustained me with their love!

Table of Contents

List of Figures

1. Introduction

The Picard group of a graph has long been an object of mathematical interest. Its motivation lies in the study of a solitaire game played on a graph called "chip-firing." This game has found applications in physics and most recently, algebraic geometry, but its combinatorial and algebraic roots prove worthy of study on their own. To construct a chip-firing game, begin by assigning an integer number of "chips" to each vertex of a graph G (this is different than simply indexing the vertex set). This establishes a game configuration on G, formally known as a divisor. At each move of the game, one chooses a vertex of G and then has the option to either "lend" or "borrow" from that vertex. In a lending move, a vertex sends one chip to each of its neighbors, boosting each neighbor's number of chips by 1 while simultaneously decreasing its own number of chips by the amount of its degree. The same happens in reverse for a borrowing move — a vertex takes one chip from each of its neighbors, which each lose a value of 1, and its own number of chips increases by the value of its degree.

Proposition 1.1. *Lending and borrowing operations are commutative, i.e. "lending then borrowing" is the same as "borrowing then lending."*

Proof. From the definition of lending and borrowing, we can record a series of chip-firing moves as a series of integer addition and subtraction operations on the free n-dimensional $\mathbb{Z}$-module where $n = |V(G)|$. This obviously commutes. $\square$

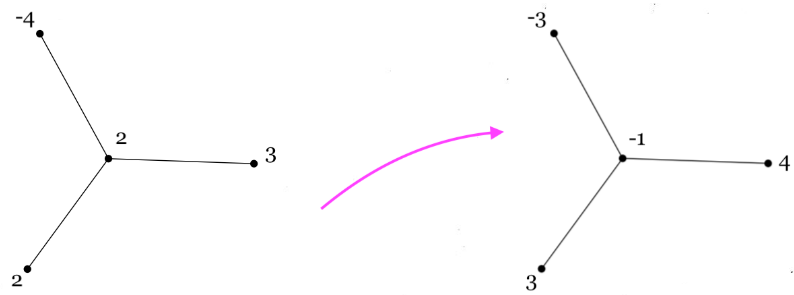

Figure 1: We lend from the central vertex.

A chip firing game is "won" when one arrives at a configuration in which each vertex has a nonnegative number of chips. There are some obvious conditions we must place on our initial configuration to ensure that winning is possible. For instance, if the total number of chips on the graph, formally known as the *degree* of the divisor, sums to less than zero, there is no way to make this happen. This is because neither lending nor borrowing affect the total number of chips — it is invariant under chip firing. As such, there will always be a vertex with a negative number of chips. The study of winnable games has largely formed the starting point for research into chip-firing. This thesis concerns a subset of that question.

Of all divisors with a net total of zero chips, the only winning configuration is the divisor in which each vertex has zero chips, which we will call the zero divisor. In any other scenario, a vertex with a positive number of chips has to be balanced by a vertex with the negative equivalent amount of chips. Thus, the only winnable degree zero configurations on a graph are somehow attained by lending and borrowing from the zero divisor. Sorting out which divisors one can reach from the zero divisor through a sequence of chip-firing moves motivates the construction of the Picard group. This is done by treating chip-firing as an equivalence relation. That is, we say that two configurations are "equivalent" if we can construct one from the other using a series of lending and borrowing moves.

Proposition 1.2. *Chip-firing is an equivalence relation.*

Proof. Checking that chip-firing satisfies the three equivalence relation axioms is straightforward.

a) A configuration is equivalent to itself — one reaches it through a sequence of zero lending and borrowing moves.

b) If we can reach configuration B from configuration A, we can retrieve configuration A by applying the same sequence of moves used to reach configuration B, except we invert lending moves to borrowing moves and vice versa.

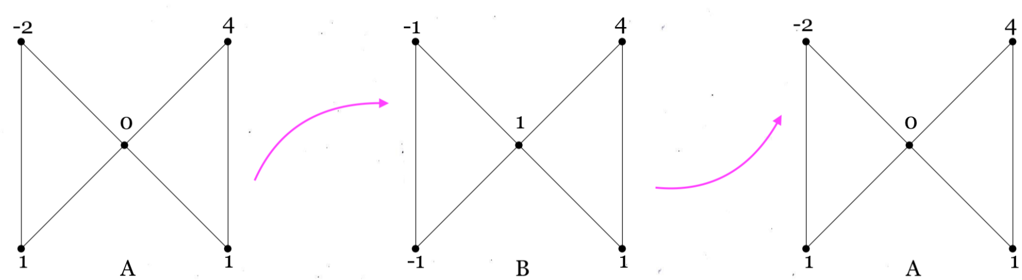

Figure 2: We lend from the bottom left vertex to produce the divisor B, and then borrow from the bottom left vertex to recover the divisor A.

c) If we can reach configuration B from configuration A, and configuration C from configuration B, we can clearly reach configuration C from A by appending the move sequence that produces C from B to the sequence that produces B from A. $\square$

The Picard group, then, is constructed by taking the set of all degree zero divisors modulo this equivalence relation. The group operation is addition where we add two configurations by summing chips on each vertex. This means the Picard group keeps track of how many distinct degree zero divisors there are, or how many degree zero configurations there are that can't be reached from each other. The Picard group encodes other valuable information about a graph G. When G is connected, the order of the Picard group of G is equivalent to the tree number of G, something we will prove later.

This thesis extends recent interests in the algebraic properties of the Picard group. In particular, we focus on the behavior of the Picard group as we "cone" over a graph G. This coning operation is performed by adjoining a single vertex to G and connecting it by a single edge to every other vertex. The n^{th} cone over G, then, is the n^{th} iteration of this coning process, or the join of G and K_n. [BMZB18] locates a set of eigenvectors of the Laplacian matrix of the n^{th} cone over G for $n \geq 1$ that correspond to divisors on the n^{th} cone. Their work associates these divisors with a subgroup of the Picard group and poses the question, when is the Picard group the direct sum of this subgroup and "everything else"? This question can also be framed as an inquiry into when a certain short exact sequence splits. As such, we inquire when the Picard group satisfies this direct-sum decomposition by asking when a graph G "splits." [GP19] picks up the work from [BMZB18] and reduces their question to the study of the Smith Normal Form of a matrix derived by modifying the original, un-coned graph Laplacian. We extend these techniques to particular graphs in order to answer questions about the frequency with which different splitting outcomes occur. We also investigate whether graph-theoretic properties of G provide information about this direct-sum decomposition. Spectral graph theory uses techniques from linear algebra to indicate graph-theoretic properties, and in some way, our thesis navigates this question in reverse.

We prove the following six theorems:

Theorem A. *If a graph G contains a universal vertex, then it splits over the n^{th} cone for all $n \geq 2$.*

Theorem B. *The 4-cycle splits only over all odd cones.*

Theorem C. *The linear graph on 4 vertices splits only over all odd cones.*

Theorem D. *The 5-cycle splits only over the n^{th} cone for $n \not\equiv 0$ mod 5.*

Theorem E. *The totally disconnected graph on k vertices splits only over the n^{th} cone when $gcd(n, k) = 1$.*

Theorem F. *Theorem F: Given a graph G, there are infinitely many $n \in \mathbb{Z}$ such that G splits over the n^{th} cone.*

Chapter 2 begins with a review of preliminary concepts from graph theory. We then formalize our notion of a divisor and the Picard group (2.1). This brings us to a discussion of the historical development of the study of chip-firing (2.2) before we outline more recent interests surrounding chip-firing's connections to algebraic and arithmetic geometry (2.3). The chapter concludes by discussing the recent "algebraic turn" in the study of chip-firing (2.4), which has focused on the group-theoretic properties of the Picard group. Section 2.4 also offers the first original insight of this thesis, in which we generalize results from [BMZB18] and [GP19] to disconnected graphs. Chapter 3 walks through the full proofs of Theorems A-F, each given their own section. We conclude by offering applications of our results and further speculations in Chapter 4.

2. Background

2.1. **Basic Definitions and Notation.** A *graph* G is a set $V(G)$ whose elements we call vertices together with a set of edges $E(G)$, 2-sets of $V(G)$ that indicate connections between vertices. The *degree* of a vertex $v \in V(G)$, denoted $deg(v)$, is the number of edges in $E(G)$ that contain v (or the number of edges *incident* to v, formally stated). We say that two vertices are *connected* in G if there is a subset of $E(G)$ that describes a path from one to the other. Two vertices are *neighbors* if they are directly linked by an edge. Finally, we say the *genus* of a graph is $g = |E(G)| - |V(G)| + 1$, as is done in [BN07]. Traditionally, this value is referred to as the *dimension* of the *cycle space* of G, but we use the term *genus* to invoke a topological analogy that we explain in Section 2.3.

There are many ways in which we can characterize or set parameters for a graph. An *undirected* graph is a graph in which edges have no orientation, whereas in a *directed* graph, edges have a "head" and "tail" chosen from the two relevant vertices. A *simple* graph is a graph in which vertices may be connected by at most one edge, and vertices may not connect to themselves. This is in contrast with a *multigraph*, in which vertices may be connected by multiple edges and a vertex can connect to itself in a loop. A *connected* graph is a graph in which all vertices are connected to each other. In a *disconnected* graph, there exists at least one pair of vertices between which no path exists. We can actually make this conception of connectedness more robust by referring to a *k-edge-connected* graph, where $k \geq 2$. We say that a graph G

is k-edge-connected if $G - W$ is connected for every set W of at most $k - 1$ edges of G. The graph on n vertices for which each vertex is connected to every other vertex is called the *complete graph* K_n. For our purposes, we will assume that G is simple and undirected, but we will not always assume that G is connected. G additionally might be a *tree*, a graph which contains no *cycles*. A cycle is a set of distinct edges that describes a path from a vertex to itself (distinctness is crucial here, as it prevents one from simply doubling back on one of the edges containing our initial vertex). A tree on n vertices has $n - 1$ edges.

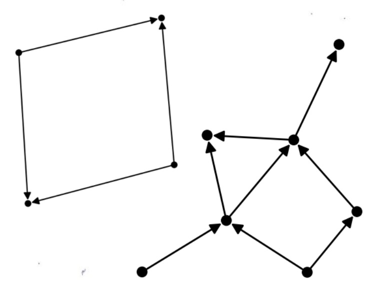

Figure 3: A simple, disconnected, directed graph

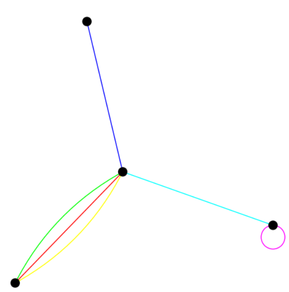

Figure 4: An undirected multigraph

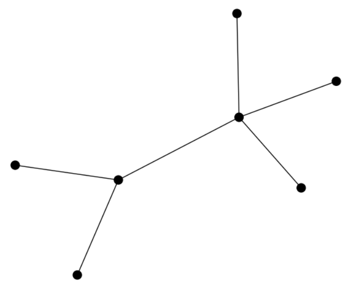

Figure 5: A tree graph

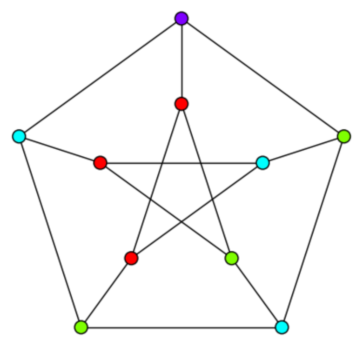

Figure 6: The Petersen graph, which is 3-edge connected

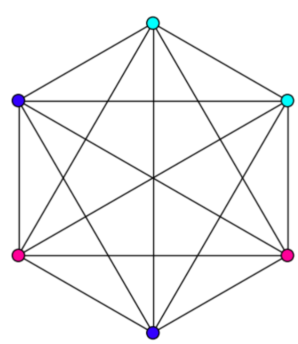

Figure 7: K_6, the complete graph on 6 vertices

There are certain matrices that we can associate to a graph. One is the *adjacency matrix*, $A(G)$. If vertices i and j are neighbors in G, then the i, j^{th} entry of $A(G)$ is the number of edges between i and j. Note that in a simple graph, the diagonal is zero and no entries are greater than 1. Another graph matrix is the *degree matrix* $D(G)$ — a diagonal matrix in which the i, i^{th} entry records the degree of the i^{th} vertex in $V(G)$. These come together to form the *Laplacian matrix* of a graph, $L(G)$, which is given by $D(G) - A(G)$. The Laplacian matrix of a graph is a symmetric, positive semidefinite matrix and it has rank $n - c$ where $n = |V(G)|$ and c is the number of connected components of G (a *connected component* of a graph G is simply a subgraph of G that is connected, but would become disconnected if we included any other vertices from $V(G)$). This means that when G is connected, the rank of $L(G)$ is $n - 1$. The kernel of $L(G)$ is spanned by the vector $(1, \ldots, 1)$ when this is the case. Additionally, $L(G)$ is diagonalizable with real eigenvalues. One graph matrix that we won't make much use of, but will be relevant in our discussion of [Big99], is the *incidence matrix* $M(G)$. This matrix is an n x m matrix where $n = |V(G)|$ and $m = |E(G)|$. Indexing the vertex and edge sets provides a basis for $M(G)$, which [Big99] defines on a directed graph. We will follow suit, as our concern with the incidence matrix does not transcend our discussion of their paper. The i, j^{th} entry of $M(G)$ is 1 if the i^{th} vertex is the "head" of the j^{th} edge, -1 if it is the "tail", and 0 otherwise. It can easily be shown that $L(G) = M^t M$.

One question we can ask about a graph is how many *spanning trees* it contains. A *spanning tree* of a graph G is a subgraph of G that is a tree which includes all vertices in $V(G)$. This question will become relevant later in the paper, as we will prove that for a connected graph, the order of the Picard group is equal to the number of spanning trees in G, also known as the *tree number* of G. As such, we provide a reproduction of Kirchoff's Matrix Tree Theorem here, which famously relates the tree number of G to the spectrum of its Laplacian matrix.

Theorem 2.1.1. *Let G be a connected, undirected graph on n vertices. Then the tree number of G is $\frac{1}{n}\lambda_1 \ldots \lambda_{n-1}$, where $\lambda_1, \ldots, \lambda_{n-1}$ are the non-zero eigenvalues of the Laplacian matrix of G.* [Sta13]

In the introduction, we motivated the study of the Picard group through the use of a chip-firing game. Here, we construct the same objects in a more formal manner. We start by defining *group*, *subgroup*, *abelian*, *normal group*, *quotient group*, and *free group* as in [DF04]. Furthermore, we refer to their definitions for *homomorphism*, *isomorphism*, *kernel*, *image*, *injective*, and *surjective*.

The set $Div(G)$ can be considered as the free abelian group on $V(G)$. We write $D(v)$ to denote the coefficient of $v \in V(G)$ in $D \in Div(G)$. Earlier, we discussed the centrality of the set of degree zero divisors in defining the Picard group. Define the degree map $deg\colon Div(G) \to \mathbb{Z}$ as the map that sends

$$\sum_{v \in V(G)} D(v)v \mapsto \sum_{v \in V(G)} D(v).$$

Then deg is a group homomorphism of $Div(G)$ to $\mathbb{Z}$ and the set of degree zero divisors is its kernel, which we denote $Div^0(G)$. We additionally referred to a notion of "won" games in our introduction, codified here as the set of *effective divisors*

$Div_+(G) = \{E \in Div(G) : E(v) \geq 0 \; for \; all \; v \in V(G)\}$. We use $Div_+^d(G)$ to denote effective divisors of degree d. Given a graph G, there is a way to construct a *canonical divisor* on G, defined as the divisor $K_G = \Sigma_{v \in V(G)}(deg(v) - 2)(v)$. This divisor has degree $2g - 2$.

If we choose an ordering $\{v_1 \dots v_n\}$ on $V(G)$ where $n = |V(G)|$, we can easily show that $Div(G)$ is isomorphic to the abelian group $\mathcal{M}(G)$ consisting of all integer-valued functions on the vertices of G as well as the space of $n \times 1$ column vectors with integer coordinates. For this reason, we will denote the column vectors corresponding to $D \in Div(G)$ and $f \in \mathcal{M}(G)$ as $[D]$ and $[f]$, respectively. This vertex ordering on $V(G)$ additionally gives a basis for $Div(G)$ as a free abelian group. We can use this basis to define the *Laplacian operator* $\Delta(G)\colon Div(G) \to Div^0(G)$, which sends

$$v \mapsto (\deg v)\, v - \sum_{wv \in E(G)} w.$$

As the name indicates, this operator has a connection to the Laplacian matrix we discussed earlier. The Laplacian matrix is the matrix representation of the Laplacian operator with respect to whatever basis we have chosen. We define $Prin(G)$ as the image of this operator Δ(G), which we call the set of *principal divisors*. Then $Prin(G)$ is a subgroup of $Div^0(G)$.

Now would be a good time to recall some of our intuition regarding chip-firing in the context of the Laplacian operator. Interestingly, the Laplacian matrix of a graph encodes information about chip-firing. This is obvious based on our definition, as the Laplacian sends the basis element $v \in V(G)$ to $deg(v)v$ and then additionally produces a $-w$ for every neighbor w of v. This describes a borrowing move on v. As such, we can model chip-firing games using the Laplacian. If we left-multiply an $n \times 1$ column vector by the Laplacian matrix, we get another n x 1 column vector. We can view this second column vector as a divisor [D] on G produced by a sequence of moves described by the left-multiplied vector V. In particular, the i^{th} coordinate of V tells us how many times we lend or borrow from v_i in our basis. Each value of 1 represents a borrowing move from v_i and each value of -1 represents a lending move. So a value of 5 in the i^{th} coordinate would indicate we borrow a total of 5 times whereas a value of -2 would indicate that we lend twice. Note that this works because lending and borrowing are commutative operations, so we only need to know the net total for how many times we lend and borrow at a vertex when describing a move sequence. Left-multiplying V by the Laplacian alone describes chip-firing from the zero divisor, but if we add their product $[D]$ to another divisor $[E]$, we can describe move sequences from that divisor. This sort of intuition motivates our next formalization.

$$\begin{bmatrix} 1 & -1 & 0 & 0 & 0 & 0 \\ -1 & 3 & -1 & 0 & 0 & -1 \\ 0 & -1 & 2 & -1 & 0 & 0 \\ 0 & 0 & -1 & 3 & -1 & -1 \\ 0 & 0 & 0 & -1 & 1 & 0 \\ 0 & -1 & 0 & -1 & 0 & 2 \end{bmatrix} \begin{bmatrix} 2 \\ 0 \\ 0 \\ 0 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 2 \\ -3 \\ 0 \\ -2 \\ 1 \\ 2 \end{bmatrix}$$

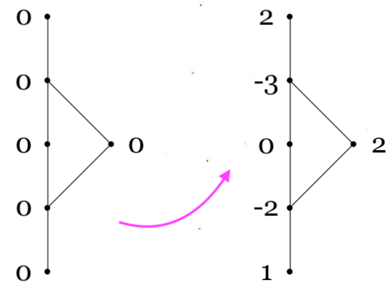

Figure 8: Here, a matrix models the outcome of borrowing twice from the top vertex on the left-bound line, once from the bottom vertex on the same line, and once from the vertex adjacent to the line.

If $D, D' \in Div(G)$, we say that D and D' are *linearly equivalent*, denoted $D \sim D'$, if $D - D' \in Prin(G)$. Since principal divisors have degree zero, linearly equivalent divisors have the same degree. Such an equivalence relation allows us to define the *linear system associated to* D, which is denoted by $|D| = \{E \in Div(G)$ *such that* $D \sim E$ *and* $E \in Div_+(G)\}$. The dimension of this linear system is $r(D)$, the *rank* of our divisor D. If $|D| = \emptyset$, $r(D) = -1$, but otherwise, for each integer $s \geq 0$, $r(D) \geq s$ if and only if $|D - E|$ is non-empty for all effective divisors E with degree s. This effectively allows us to keep track of how "winnable" a given game configuration is, as subtracting effective divisors from D essentially removes a nonnegative number of chips from each vertex. The rank of a divisor, then, lets us know how many chips we can arbitrarily take from a divisor before it is no longer winnable.

Lemma 2.1.2. *For all $D, D' \in Div(G)$ such that $r(D), r(D') \geq 0$, we have that $r(D + D') \geq r(D) + r(D')$.* [BN07, Lemma 2.1]

Proof. Let E' be an arbitrary effective divisor of degree $r(D')$, let E be an arbitrary effective divisor of degree $r(D)$, and let $E'' = E + E'$. Then $|D - E|$ and $|D - E'|$ are both non-empty, meaning $D - E \sim F$ and $D - E' \sim F'$, where F and F' are both effective. By assumption, $(D + D') - (E + E') = (D + D') - E'' \sim F + F'$. Since F and F' are both effective, $F + F'$ is effective, which implies $(D + D') - E''$ is winnable. From how we have constructed E, E', and E'', it follows that $r(D + D') \geq r(D) + r(D')$. $\square$

This notion of linear equivalence is in accordance with our chip-firing relation posited in Section 1.1. If $D - D'$ is in $Prin(G)$, the image of $\Delta(G)$, then there is an $n \times 1$ column vector V such that $[D'] + L(G) \times V = [D]$, i.e. a move sequence that brings us from D' to D as per our explanation above. With all of this in mind, we are naturally motivated to define the Picard group of a graph G.

Definition 2.1.3. The Picard group of a graph G is the quotient group $Pic^0(G) = Div^0(G)/Prin(G)$.

We can represent the equivalence class of a degree zero divisor D in $Pic^0(G)$ as (D), but often we will refer to it just the same as $D \in Pic^0(G)$. This just about establishes the context necessary to move forward, but we can ask, are there any "nice" ways to represent these equivalence classes? This would provide a "nice" way to list elements of $Pic^0(G)$. It turns out that there is such a method. Define a *q-reduced divisor* on G as follows: Given a graph G, arbitrarily pick a vertex and label it "q." A q-reduced divisor D is such that (i) $D(v) \geq 0$ for all $v \neq q$ and (ii) for every non-empty set $A \subset V(G)$ with $q \notin A$, lending from each element of A at the same time produces a configuration where $D(v) < 0$ for some $v \in A$.

Lemma 2.1.4. *A divisor D is q-reduced if and only if D is effective outside q, but for every non-constant function $f \in \mathcal{M}(G)$ having a global maximum at q, the divisor corresponding to [D] + $L(G) \times [f]$ is not effective outside q.* [Bak, Lemma 1.1]

Proof. We will only rely on the "if" direction of this statement, so we only reproduce this portion of the proof. If D is q-reduced and $f \in \mathcal{M}(G)$ is a non-constant function with a global maximum at q, then let $S \subset V(G)$ be the set of vertices where f attains its minimal value. Since f has a maximum at q and f is non-constant, S is non-empty and $q \notin S$. Since D is q-reduced, there is $v \in S$ such that $D(v) < outdeg_S(v)$, where $outdeg_S(v)$ measures the number of edges that connect v outside of the set S. For all $x \in S$, $E(x) \leq -outdeg_S(x)$ where E is the divisor $L(G) \times [f]$. This is because f has a distinct maximum and minimum by virtue of being non-constant, so $E(x)$ is negative. Furthermore, f models at least one lending move from each $x \in S$. At best, any chips lent to other vertices in S are replenished, i.e. lent back. So $E(x)$ loses a net total at least the amount $outdeg_S(x)$. Then for v as before, $(D + L(G) \times f)(v) < 0$ so $D + L(G) \times f$ is not effective outside q. □

Proposition 2.1.5. *Each equivalence class of $Pic^0(G)$ contains a unique q-reduced divisor.* [BN07, Proposition 3.1]

Proof. We first show existence. If given a divisor D, order $V(G)$ $(v_1 \dots v_n)$ with n

$= |V(G)|$ and $v_1 = q$ such that each vertex except q has a neighbor that precedes it in the order. We can make $D(v) \geq 0$ for all $v \neq q$ by working step by step from the end to the front of this order. At each step i, locate a neighbor of v_{n+1-i} that precedes it in the order. Lend from this neighbor until $D(v_{n+1-i}) \geq 0$ (Since we don't change the equivalence class of D by doing this, we will continue to call it D even though the configuration technically changes.). At the end of this process, D satisfies our first condition. Now, either D is q-reduced, or there is a nonempty set $A \subset V(G)$ such that $q \notin A$ and we can lend from each $v \in A$ simultaneously without any v-coefficients dropping below zero. If the latter is the case, lend from each $v \in A$ repetitively. Obviously, if we keep doing this, at some point we will get $D(v) \leq 0$ for some $v \in A$. So, repeat until you get a q-reduced divisor.

Now we claim that this q-reduced divisor is unique. Suppose D and D' are q-reduced such that $D \sim D'$ with $D \neq D'$. Then $[D - D'] = L(G) \times [f]$ for some function $f \in \mathcal{M}(G)$. We know f is non-constant, and since D' is effective outside q and D is q-reduced, it follows from Lemma 2.1.7 that f does not have a maximum at q. So the set S of vertices where f attains a maximum is a subset of $V(G)$ which does not contain q. We claim that $E(x) \geq outdeg_S(x)$ for all $x \in S$ where E is the divisor $L(G) \times [f]$ by the same argument as in Lemma 2.1.7. Since D' is q-reduced, there is some $v \in S$ such that $D'(v) < outdeg_S(v)$. Since $D(v) \geq 0$, we get $D'(v) < outdeg_S(v) \geq E(v) \geq (D+E)(v) = D'(v)$. This is a contradiction, so D' can not also be q-reduced. $\square$

The notion of a q-reduced divisor will be most relevant in our discussion of recent applications of chip-firing to algebraic geometry. In [BN07], the authors use q-reduced divisors en route to a graph-theoretic analogue of the Riemann–Roch criterion. In practice, q-reduced divisors are handy in checking relationships between divisors, such as those of the form $nD \sim D'$ where $n \in \mathbb{Z}$, as we can q-reduce either side and immediately see if the two are equivalent.

2.2. **Combinatorial Foundations for Chip Firing.** In this section, we take an aside to discuss two variants on our chip-firing game explicated in [BLS91] and [Big99]. These constructions in fact comprise the original constructions of chip-firing and are thus worthy of our attention. Indeed, these papers have functioned as introductions to chip-firing and can provide insight into the ways in which the field was previously self-motivated. The blend of combinatorics and algebra in these surveys sets the stage for the conceptual tools we use now and is thus instructive in providing historical context for our results. These papers also provide an important genealogy for the terminology used in this paper.

In [BLS91], the authors approach chip-firing with a heavily combinatorial eye. They trace the lineage of their results to a handful of different "balancing games," which they depart from through the imposition of additional rules and in their explicitly setting the game on a graph. Their construction of chip-firing is as follows: Put a pile of chips on each vertex of a finite, connected, simple graph G, summing to N chips altogether. Instead of using both lending and borrowing moves to alter G, we are restricted only to lending. Furthermore, moves are constricted in a manner unlike our iteration of the game. In order to lend (or, in their terms, "fire") from a vertex, it needs to have at least as many chips as its degree. The game terminates, or is "won," when each vertex has a number of chips less than its own degree. When a game can be won, we say that it is *finite*.

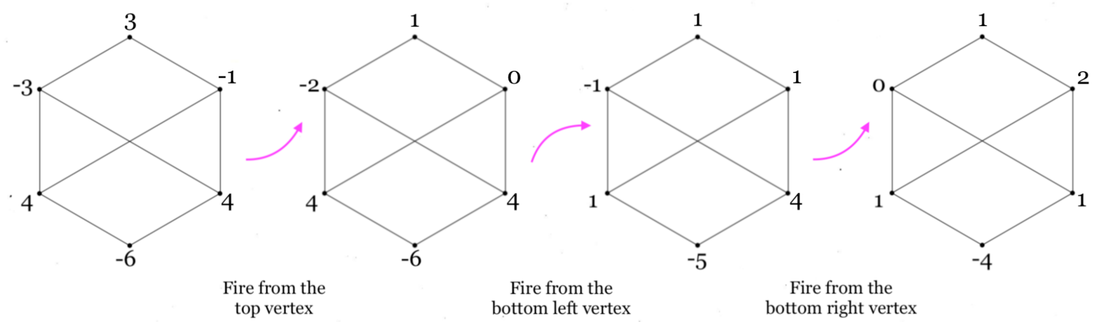

Figure 9: The demonstration of a finite game.

Lemma 2.2.1. *If a chip-firing game never terminates, then every vertex is fired infinitely often.* [BLS91, Lemma 2.1]

Lemma 2.2.2. *If a chip-firing game terminates, then there is a vertex which is never fired.* [BLS91, Lemma 2.2]

Proof. We show that if every vertex on a graph is fired, then the game cannot terminate. This is because if we fire all but one vertex v, each of v's neighbors has fired. As such, v has received a chip from each of its neighbors and now has at least $deg(v)$ chips. This means we can fire v. Once we've done this, the first vertex we fired on the first go around has experienced every other vertex firing but itself, and similarly must be able to fire as well. Since this process always produces a vertex that is the "longest-waiting" vertex, there is always a vertex to fire and our initial game cannot terminate. □

The authors concern themselves with the finiteness of different chip-firing games. They articulate that finiteness depends on choice of graph and starting configuration, but not the choices one makes during the course of a game. Really, it is only the choice of graph and starting configuration that determine both the terminating configuration and the number of steps necessary to attain it. [BLS91] is interested in offering basic heuristics to indicate when a game is winnable as well as bounds on the maximum number of steps one must take to go from an initial to terminating configuration. Along the way, they heavily rely on language theory and make some use of the graph Laplacian.

As their chip-firing only requires lending, [BLS91] records moves as a "word" string of vertex "letters" that indicates the sequence of which vertices we fire from. Formally, they take a set E and define a *language* L over E as a set of finite strings formed from elements of E. If we have a word $a \in L$, we take [a] to be the *score* of that word, a vector in $\mathbb{Z}^n$ ($n = |E|$) where the i^{th} coordinate records the number of times the i^{th} letter of E appears in a. A language is *left-hereditary* if whenever it contains a string, it also contains every beginning section of that string. For instance, English is not left-hereditary because even though "chip" is a word, "c," "ch," and "chi" are not. A language L over a set E is *locally-free* if for a word $a \in L$ and $x, y \in E$ such that $x \neq y$, if $ax \in L$ and $ay \in L$, then $axy \in L$. We also say L is *permutable* if whenever $a, b \in L$, $[a] = [b]$, and $ax \in L$ for some $x \in E$, $bx \in L$. Abstractly, [BLS91] is concerned with when languages possess the "strong exchange" property, which states that if $a, b \in L$, then a contains a subword a' such that $ba' \in L$ and $[ba']$ is the coordinate-wise maximum of $[a]$ and $[b]$. The first lemma that [BLS91] proves is as such.

Lemma 2.2.3. *Every locally free permutable left-hereditary language has the strong exchange property. Conversely, every language with the strong exchange property is locally free and permutable.* [BLS91, Lemma 1.2]

As it goes, the language which is the set of legal games on a given graph G endowed with an initial configuration D is a locally free permutable left-hereditary language. [BLS91] additionally sets out to define a *basic* word, which is a word that is not a *prefix* or *initial segment* of any other word in the language. If a left-hereditary language with the strong exchange property has any basic words, all basic words have the same length and there are no words that are any longer. Otherwise, we could append strings to basic words by the strong exchange property, contradicting their definition. The common length of basic words is referred to as the *rank* of a language (this rank is infinite if there are no basic words).

[BLS91] defines an equivalence relation on words in which two words $a, b \in L$ are equivalent if for every string of letters c, $ac \in L$ if and only if $bc \in L$. They use the equivalence classes of this relation, which they call *flats*, combined with the previous result to show that in a language of legal games on G starting from D with finite rank, two legal games lead to the same position if and only if they have the same score. Finite games, games that terminate, that is, make for languages that have finite rank as terminating legal games constitute basic words. This final result indicates, then, that all terminating legal games have the same length and score and necessarily lead to the same configuration. This outlines a proof for one of their primary results.

Theorem 2.2.4. *Given a connected graph and an initial distribution of chips, either every legal game can be continued indefinitely, or every legal game terminates after the same number of moves with the same final position. The number of times a given vertex is fired is the same in every legal game.* [BLS91, Theorem 1.1]

The rest of the paper walks through the two remaining important results. The first concerns the relationship between N, the total number of chips in a game, and the game's finiteness.

Theorem 2.2.5. *Let G be a connected graph on n vertices and m edges with N chips. Then, if $N > 2m - n$, the game is infinite. If $m \leq N \leq 2m - n$, there exists an initial configuration guaranteeing finite termination and also one guaranteeing an infinite game. If $N < m$, then the game is finite.* [BLS91, Theorem 2.3]

We will not reproduce a proof of this result here. However, note that some of these results are quite intuitive. If we put too many chips on a graph, there is no way to allocate them such that each vertex has fewer chips than its degree. When $N > 2m - n$, the most diffuse distribution of chips still places at least $deg(v) - 1$ chips at each vertex with certain vertices necessarily carrying more (thus, being able to fire). If we don't put enough chips on the graph, then necessarily we can get to a point where each divisor has fewer chips than its degree. The question remains, if we know how to categorize finite versus infinite games, and we know that all legal terminating games on a winnable configuration are equivalent, is there any way to estimate just how many moves a legal terminating game will entail? This is the final problem that [BLS91] tackles.

Theorem 2.2.6. *Let G be a connected graph on n vertices where λ_1 is the smallest non-zero eigenvalue of the Laplacian matrix. Then the number of steps in any terminating chip-firing game with N chips is at most $2nN/\lambda_1$.* [BLS91, Lemma 3.2]

We refer the interested reader to [BLS91] for a proof of this result, but note the significance in the presence of λ_1. The second smallest eigenvalue of a graph Laplacian is known to encode information about the connectivity of the graph. As we will see later, this question of connectivity is relevant to the construction of configurations and more importantly, the structure of the Picard group.

[Big99] constructs a variant of the chip-firing game that is largely identical to the one described in [BLS91]. We assume that G is a finite, connected multigraph without loops. Biggs assumes that G has an orientation, but this is largely a technical assumption that allows for certain constructions and does not pertain to the paper's major results. Once again, we assign an integer number of chips to each vertex and allow a vertex to fire when its number of chips is at least its degree. However, this time around, we arbitrarily pick some vertex $q \in V(G)$ to perpetually hold the negative sum of chips from the rest of the vertices in $V(G)$. This way, the total number of chips in the game is zero (all other vertices must have at least zero chips). Whereas the other vertices fire under the same conditions as before, q fires if and only if no other firing is possible, i.e. when the game terminates. Biggs refers to configurations where vertices besides q have fewer chips than their degree as *stable*. A sequence of firings is *q-legal* when vertices besides q only fire when they have more chips than their degree and q only fires after the game reaches a stable configuration. A firing sequence is *proper* if it does not contain a q-firing. We say that a configuration D is *recurrent* when there is a q-legal sequence from D such that D recurs upon its completion. A configuration is *critical* when it is both stable and recurrent.

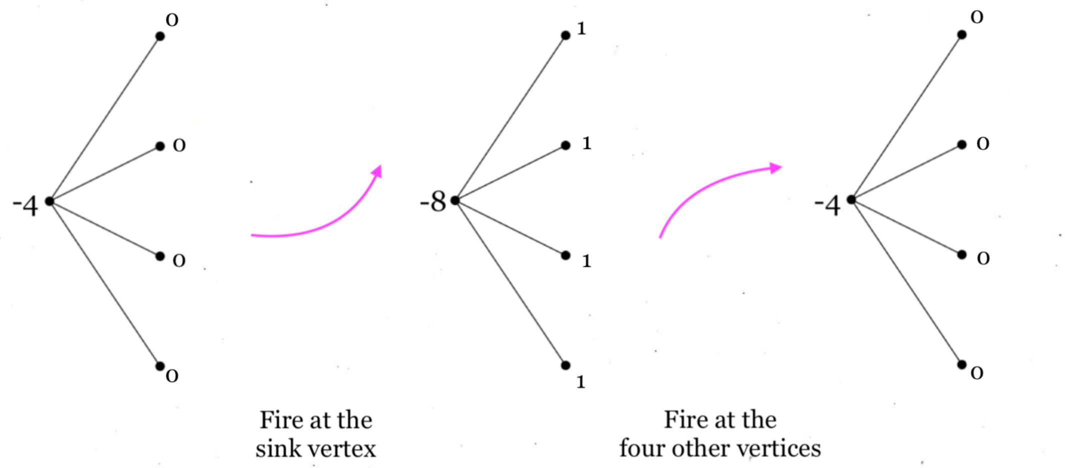

Figure 10: A demonstration of a critical configuration recurring

Proposition 2.2.7. *Given a configuration D on a graph G, there is a q-legal firing sequence that produces a critical configuration.* [Big99, Lemma 3.2]

Proof. We can use a proper firing sequence to arrive at a stable configuration. If we fire from q, we can use another proper firing sequence to get back to a stable configuration. If we do this infinitely many times, there are only finitely many stable configurations on a graph G, so necessarily, one of these configurations recurs. □

[Big99] shows this critical configuration is unique by constructing a "nice" way to combine proper firing sequences. This method allows them to in effect "triangulate" stable configurations. If one proper firing sequence X leads to a stable configuration, there is a way to append a modified version of X to any other proper firing sequence such that the combination produces the same stable configuration ([Big99], Corollary

3.5). With this construction, they show that the set of critical configuration equivalence classes is well-defined. Their next challenge is to endow this set with a group structure.

Instead of explicitly giving the set of critical configurations, which they denote $\kappa(G)$ and call the *critical group*, a group operation, they prove it is in bijection with another abelian group. Consider the map $\sigma\colon \mathcal{M}(G) \to \mathbb{Z}$, which sends integer-valued functions f to the sum $\Sigma_{v \in V(G)} f(v)$ (this is the same map as the degree map we defined earlier).

Proposition 2.2.8. *The image of the graph Laplacian is a normal subgroup of* $Ker(\sigma)$*.* [Big99, Lemma 4.1]

Proof. The product σM is the zero matrix since M has only two non-zero entries in each column and one is 1 and the other is -1. If $x \in Image(L(G))$, there is y such that $x = L(G)(y) = MM^t(y)$. Then $\sigma(x) = \sigma(MM^t(y)) = (\sigma M)(M^t(y)) = 0$. The image of the Laplacian is thus a subgroup of $Ker(\sigma)$ and since $Ker(\sigma)$ is abelian, it is a normal subgroup. □

Proposition 2.2.9. $\kappa(G)$ *is isomorphic to* $Ker(\sigma)/Image(L(G))$*.* [Big99, Theorem 4.2]

Proof. First, every equivalence class $(f) \in Ker(\sigma)/Image(L(G))$ corresponds to a configuration on G. Given $f \in Ker(\sigma)$, let l be the configuration given by $l(u) = deg(u) - 1$ when $f(u) \geq 0$ and $l(u) = deg(u) - 1 - f(u)$ when $f(u) < 0$ for $u \neq q$.

Then $l(q) = -\Sigma_{u \neq q} l(u)$ as always. There is a finite firing sequence that reduces l to a stable configuration k. Represent this firing sequence by the vector v, then $k = l + L(G)v$. Let $z = f + k - l$. Then $z = f - L(G)(v)$, so $(z) = (f)$ and $z(u) = f(u) + l(u) - k(u) \geq deg(u) - 1 - k(u) \geq 0$. Thus, we can use the configuration z to represent (f).

Define the map $h\colon Ker(\sigma)/Image(L(G)) \to \kappa(G)$ by $h(\alpha) = \gamma(b)$, where b is any one of the configurations corresponding to some $(\alpha) \in Ker(\sigma)/Image(L(G))$ and $\gamma(b)$ is its unique linearly equivalent critical configuration. h is a surjection since given $c \in \kappa(G)$ we can consider c as an element of $\mathcal{M}(G)$ and then $h(c) = \gamma(c) = c$. If $h(s_1) = h(s_2)$, then $\gamma(s_1) = \gamma(s_2) = c$. This implies there are vectors x_1, x_2 such that $s_1 + L(G)(x_1) = c$ and $s_2 + L(G)(x_2) = c$, so $s_1 + L(G)(x_1 - x_2) = s_2$ and $(s_1) = (s_2)$. Thus, h is also injective, so it is an isomorphism. $\square$

Note that this is how we have constructed our Picard group. This fact will soon be relevant, as the rest of the paper sets up isomorphisms between $\kappa(G)$ and various other groups associated to a graph. By the end of their exposition, they use these isomorphisms to show that the order of $\kappa(G)$ is the tree number of the graph.

First, they use the theory of flows and cuts to construct a lattice defined from the incidence matrix. Denote $C^0(G;\mathbb{Z})$ as the vector space of real-valued functions on $V(G)$ and $C^1(G;\mathbb{Z})$ as the vector space of real-valued functions on $E(G)$. We will simply refer to these as C^0 and C^1 for short. We can define an inner product relation on C^1:

$$\langle x, y\rangle = \sum_{e\in E(G)} x(e)y(e).$$

[Big99] indicates that the incidence matrix and its transpose $M^t(G)$ are homomorphisms $M\colon C^1 \to C^0$ and $M^t\colon C^0 \to C^1$. As per linear algebra, C^1 is the direct sum of $Ker(M)$ and its orthogonal complement, here denoted by Z and B, respectively. Next, [Big99] uses C_I to denote integer-valued functions on $E(G)$, which considered as a subset of C^1 is a lattice.

Definition 2.2.10. Given a field Q, an n-dimension vector space V over Q with basis $\{v_1, \ldots, v_n\}$, and a ring $R \subset Q$, the *lattice* L defined by R is $L = \{\Sigma_{i=1}^n a_i v_i | a_i \in R\}$.

[Big99] considers the C_I version of the direct sum $Z \oplus B$ they constructed on C^1. First, define $Z_I = Z \cap C_I$ and $B_I = B \cap C_I$. Then $Z_I \oplus B_I$ is a proper sub-lattice of C_I, though not an isomorphic copy ([Big99], Theorem 6.1). This is important, because it indicates that $C_I/Z_I \oplus B_I$ is nontrivial and in particular, they show its order is the tree number of G. Their argument uses the construction of various isomorphisms, the most important of which is from $C_I/Z_1 \oplus B_1$ to the group $B_I^\sharp/B_I$, where $B_I^\sharp = \{x \in B\colon \langle x, b\rangle \in \mathbb{Z} \text{ for all } b \in B_I\}$ is the *dual* of B_I. Interestingly enough, they also refer to an isomorphism between $C_I/Z_1 \oplus B_1$ and $Z_I^\sharp/Z_I$, the latter of which is the Jacobian group. As shown in [BN07], this Jacobian is isomorphic to our Picard group.

Our discussion of cuts and flows culminates with this result on the order of $C_I/Z_I \oplus B_I$. The paper carries on with a discussion of the Picard group, which

[Big99] articulates is rooted in algebraic geometry. Formally, this Picard group is defined here as a quotient $M(C_I)/M(D_I)$, where M is once again the incidence matrix. $M(C_I)$ is referred to as the group of divisors of degree zero and $M(D_I)$ is the group of principal divisors, though their names are meant to refer to geometric constructs, not our chip-firing objects (the analogous language will find explanation in our next section). [Big99] shows that $M(C_I)$ is equivalent to $Ker(\sigma)$ and D_I is equivalent to the image of M^t. Since $\mathrm{MM}^t = \mathrm{L(G)}$, this indicates that $M(D_I)$ is equivalent to the image of our Laplacian. Conveniently, this proves that the Picard group is isomorphic to $\kappa(G)$.

At this point, we've almost drawn enough isomorphisms to get what we need. [Big99] refers to a result from [Big97] which establishes an isomorphism between the Picard group and $B_I^\sharp/B_I$. As $B_I^\sharp/B_I$ is isomorphic to $C_I/Z_1 \oplus B_1$, so are the Picard group and $\kappa(G)$. As we indicated earlier that the order of $C_I/Z_1 \oplus B_1$ is the tree number of G, this implies that the order of $\kappa(G)$ is also the tree number.

[Big99] concludes with a handful of observations that are theoretically relevant and instructive in the practice of determining the Picard group of particular graphs. The paper refers to a *reduced* Laplacian matrix, which differs from the regular graph Laplacian by the deletion of one row and column. We denote the reduced Laplacian by $\tilde{L}(G)$. When G is connected, this procedure is useful because it allows us to work with a matrix of full rank. As [Big99] indicates, this reduced Laplacian is the matrix of relations for the Picard group, which is isomorphic to the cokernel of $\tilde{L}(G)$. This is the case irrespective of our choice of row and column to delete. It is somewhat conventional in the literature on chip-firing to work with the reduced Laplacian for simplicity's sake, as is done in [BMZB18] and [GP19]. Since the Picard group is a finite abelian group, by the Fundamental Theorem of Finite Abelian Groups, it has a direct-sum decomposition into cyclic invariant factors. One way to indicate these invariant factors is by taking the Smith Normal Form of our graph Laplacian. This is functionally what we will do to tackle the problems outlined in 2.4. The last two sections of [Big99] locate critical configurations of a certain order on different families of graphs as a means of identifying invariant factors of the Picard group. This will be highly relevant to our course of action, as this is how the direct-sum decomposition in [BMZB18] is posed.

2.3. **The Geometry of a Divisor.** [Big99] alludes to an analogy between $\kappa(G)$ and a motivation for our Picard group that originates in algebraic geometry. In the past decade or so, researchers have zeroed in on a series of tools that allows for the transference of results from algebraic geometry to graph theory via chip-firing, and vice versa. This trend has even attracted journalistic coverage (e.g. [Har18]) as of late. Section 2.3 will give an overview of two papers that work in this tradition, one of which most notably provides a graph-theoretic analogue for the Riemann–Roch Theorem, a classical result in algebraic geometry that relates the topological and analytic properties of a Riemann surface. We will briefly explicate the exchange between these two fields before discussing the implications for chip-firing more directly. First, it would help to set out our definitions.

We use the definition of *ring* provided in [DF04] and assume that all rings are *commutative with unity*. [DF04] furthermore provides our definitions for *quotient field*, *algebraic closure*, *residue field*, and *discrete valuation ring*. We will often refer to a *complete* discrete valuation ring, which is a discrete valuation ring R endowed with a topology such that R is complete under that topology, i.e. every Cauchy sequence in R converges in R. A good deal of our algebraic geometric work is based in the theory of schemes, which will not be explicated here. The interested reader can refer to texts such as [Har77], [Nel], and [Vak]. The theory of divisors in algebraic geometry would obviously relate to a study of curves, so we use the notion of a *smooth curve* as is found in [Bak08]. We also import from [Bak08] the notion of a *strongly semistable regular model* $\mathfrak{X}$ of a curve X and its *special fiber* $\mathfrak{X}_k$, where X is defined over the quotient field of a complete discrete valuation ring. Our interest will be in the irreducible components of this special fiber $\mathfrak{X}_k$, where the meaning of *irreducible* is taken in the scheme-theoretic sense.

Earlier, we defined the group of divisors on a graph. We can similarly define a divisor on a scheme, and in fact, this notion of a divisor is where our graph-theoretic construction originates. Let X be an integral, locally Noetherian scheme. There are two notions of a divisor that we can define over X — a Weil divisor and a Cartier divisor. When X is regular, these are equivalent. Since our work only concerns regular schemes, we will stick to defining Weil divisors, which are more intuitive. A *prime* *divisor* on X is an integral, closed sub-scheme $P \subset X$ where $dim(X) - dim(P) = 1$. The group of *Weil divisors* $Div(X)$, then, is the free abelian group on the prime divisors of X, where elements are $\Sigma_P\ a_P P$ with only finitely many $a_P \neq 0$. We say that a divisor $D \in Div(X)$ is *effective* when $a_P \geq 0$ for all P. Given a non-zero rational function f on X, there is a natural way to associate f with a divisor, equal to $\Sigma_P\ ord_P(f)P$, where $ord_P(f)$ is the length of the longest chain of submodules in $\mathcal{O}_{X,P}/(f)$. We say that a divisor in $Div(X)$ is *principal* when it is the divisor of a non-zero rational function on X. The group of principal divisors, $Prin(X)$, is a subgroup of $Div^0(X)$, just as in the graph-theoretic case.

One arrives in the graph-theoretic world of chip-firing from the algebraic/geometric world of curves via a simplification process broadly known as *tropicalization*. To tropicalize a curve, one sets out its fundamental properties and finds a way to reconstruct them in a simpler fashion. In particular, [Bak08] uses the notion of a *dual graph* in which vertices correspond to the irreducible components of a curve and edges connect vertices where irreducible components intersect. Let R be a complete discrete valuation ring, let K be its field of fractions, and let k be its algebraically closed residue field. As indicated above, given a smooth curve X over K, we can define $\mathfrak{X}$ over R, the strongly semistable regular model of X with special fiber $\mathfrak{X}_k$. Let $\mathfrak{C} = \{C_1, ..., C_n\}$ be the set of irreducible components of $\mathfrak{X}_k$. Then, unless noted otherwise, in this section we will use G to denote the dual graph of some special fiber $\mathfrak{X}_k$. Since $\mathfrak{X}$ is strongly semistable, elements of $\mathfrak{C}$ are smooth and do not self-intersect, so G is well-defined and contains no loops.

[Bak08] constructs a *specialization map* $\rho\colon Div(\mathfrak{X}) \to Div(G)$, a homomorphism that sends a divisor D on the model $\mathfrak{X}$ to a divisor $\bar{D}$ on the dual graph G. This homomorphism embeds the principal divisors of $\mathfrak{X}$, $Prin(\mathfrak{X})$, as a subgroup of $Ker(\rho)$. The formula for this map is as follows:

$$\rho(D) = \Sigma_{v_i}(C_i \cdot D)(v_i)$$

where $(C_i \cdot D) = deg(\mathcal{O}_{\mathfrak{X}}(D)|_{C_i})$. [Bak08] indicates a means by which we can send divisors $D \in Div(X)$ to $Div(\mathfrak{X})$. This allows us to think of the same map ρ as a map from $Div(X)$ to $Div(G)$. When $\rho\colon Div(X) \to Div(G)$, it has the following properties: (i) ρ is a degree preserving homomorphism, (ii) if $D \in Div(X)$ is effective, then $\rho(D) \in Div(G)$ is effective, (iii) if $D \in Prin(X)$, the $\rho(D) \in Prin(G)$.

When $D \in Div(X)$, we call $\bar{D} = \rho(D) \in Div(G)$ its *specialization to* G (or its *specialization*, more simply). We can ask how various properties of D change when we move to $\bar{D}$. For instance, we indicated above how properties such as the degree, effectiveness, and principality of D are invariant under specialization. [Bak08] takes interest in the behavior of $r(D)$. We have a notion of $r(D)$ when D is a divisor on a graph G, but we must define what $r(D)$ looks like when D is a divisor on a curve X. When it is not immediately apparent which of the two we are discussing, we will use $r_G(D)$ or $r_X(D)$.

Definition 2.3.1. Let X be a smooth curve over a field K. Then for $D \in Div(X)$, $|D| = \{E \in Div(X)\colon E \text{ is effective and } E \sim D\}$. We say $r(D) = -1$ when $|D| = \emptyset$.

Otherwise, $r(D) = \max\{k \in \mathbb{Z} \mid |D - E| \neq \emptyset \text{ for all } E \in Div_+^k(X)\}$.

Lemma 2.3.2. *(Specialization Lemma) For all $D \in Div(X)$, we have $r_G(\rho(D)) \geq r_X(D)$.* [Bak08, Lemma 2.8]

So, the rank of a divisor does not decrease under specialization. [Bak08] uses this to furnish a handful of different characterizations of graphs and smooth curves. Most notably, they relate the notion of a Weierstrass point on an algebraic curve to an equivalent notion of a Weierstrass point on a graph. Classically, a Weierstrass point refers to a point on a curve X at which a non-constant rational function on X has a pole of order no larger than the curve's genus. This function additionally has poles at no other point on the curve. [Bak08] defines a Weierstrass point on a graph as a vertex $v \in V(G)$ such that the divisor $g(v)$, where g is the genus of G, has rank at least 1. While these definitions accord in many important ways, their meanings diverge slightly. For instance, while a curve with genus 0 or 1 has no Weierstrass points and any curves with $g \geq 2$ necessarily do, there exist graphs with $g \geq 2$ that do not. What is most important, as we will see, is that this notion of a Weierstrass point on a graph is strong enough to retain information about Weierstrass points on curves under specialization. Along with a Weierstrass point, we define a Weierstrass gap. Algebraically, we define a Weierstrass gap at a point p on a curve as an integer m such that no rational function defined on the curve has an m-order pole exclusively at p. For a graph with $p \in V(G)$, a Weierstrass gap at p is an integer $k \geq 1$ such that the rank of the divisor $k(p)$ is equal to the rank of the divisor $(k-1)(p)$. There are a few properties of graph Weierstrass points that [Bak08] lays out plainly.

Lemma 2.3.3. *The following are equivalent:*

(i) P is a Weierstrass point.

(ii) There exists a positive integer $k \leq g$ which is a Weierstrass gap at P.

(iii) The rank of the canonical divisor minus $g(P)$ is greater than zero.

Additionally, the set of Weierstrass gaps at a point $P \in V(G)$ is contained in $\{1, 2, \ldots, 2g-1\}$ and has cardinality g. [Bak08, Lemma 4.2]

Lemma 2.3.4. *Let v be a vertex of a graph G of genus $g \geq 2$, and let G' be the graph obtained by deleting the vertex v and all edges incident to v. If G' is a tree, then v is not a Weierstrass point.* [Bak08, Lemma 4.7]

An arithmetic surface $\mathfrak{X}$ over a ring R is *totally degenerate* if the genus of its dual graph G is the same as the genus of X, the curve it models. [Bak08] uses the Specialization Lemma as a way to associate Weierstrass points of X with Weierstrass points of G, and vice versa. Since G is much more intuitive to work with, the results provided in [Bak08] allow us to detect Weierstrass points on X in a much simpler fashion. To start, we can impose conditions on $\mathfrak{X}$ to ensure that Weierstrass points are preserved under the Specialization Map.

Theorem 2.3.5. *If $\mathfrak{X}$ is a strongly semistable, regular, and totally degenerate arithmetic surface, then for every K-rational Weierstrass point $P \in X$, $\rho(P)$ is a Weierstrass point of the dual graph G of $\mathfrak{X}$.* [Bak08, Corollary 4.9]

One could predict a result like this based on the trajectory of this section thus far — impose conditions on $\mathfrak{X}$ until one finds that certain properties are preserved under the specialization map. The next result more widely proves how we can better understand the original curve X based on its dual graph G.

Theorem 2.3.6. *(a) Let $\mathfrak{X}/R$ be a strongly semistable, regular, totally degenerate arithmetic surface whose special fiber has a dual graph with no Weierstrass points. Then X does not possess any K-rational Weierstrass points.*
(b) Let $\mathfrak{X}/R$ be an arithmetic surface whose special fiber consists of two genus 0 curves intersecting transversely at 3 or more points. Then every K-rational Weierstrass point of X specializes to a singular point of $\mathfrak{X}_k$.
(c) More generally, let $\mathfrak{X}/R$ be a strongly semistable and totally degenerate arithmetic surface whose dual graph G contains a vertex v for which $G' := G\backslash\{v\}$ is a tree. Then there are no K-rational Weierstrass points on X specializing to the component C of $\mathfrak{X}_k$ corresponding to v. [Bak08, Corollary 4.10]

These results hopefully demonstrate the power of chip-firing in ascertaining information about algebraic curves. The theorems above should explain the interest that algebraic geometers have taken in the field and in the use of chip-firing's algebraic and graph-theoretic machinery. As a bonus, [BN07] is able to import classic theorems from algebraic geometry into the realm of chip-firing. We specifically get a graph-theoretic analogue of the foundational Riemann–Roch Theorem.

Theorem 2.3.7. *(Riemann–Roch Theorem for Graphs) Let G be a graph and D a* *divisor on G. Then $r(D) - r(K_G - D) = deg(D) + 1 - g$.* [BN07, Theorem 1.12]

We can informally think of genus as encoding the number of holes an object has, both in the context of surfaces and graphs. The classical Riemann–Roch Theorem is profound, then, because it relates the number of holes on a surface, which is a geometric or topological invariant, to the rank of a divisor, which relates more directly to algebraic properties of the surface by way of the canonical divisor, but which is also contingent upon choice of divisor. On a graph, the genus measures a similar invariant by keeping track of how many more edges than vertices there are. As this number increases, we see a similar increase in the number of cycles on the graph, which we can think of similarly of holes. Our Riemann–Roch Theorem for Graphs, then, relates this graph invariant to the "winability" of different game configurations.

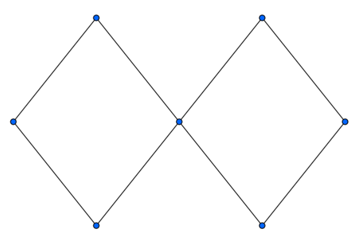

Figure 11: A genus 2 graph

There are a few immediate consequences of the theorem. The first is that the canonical divisor K_G of a graph G has rank $g-1$, since it has degree $2g-2$, $K_G - K_G$ is the zero divisor, and the rank of the zero divisor is zero. Second, when D has degree $g-1$, then it is winnable if and only if the divisor $K_G - D$ is winnable, since the two necessarily have the same rank. Finally, when $deg(D)$ is greater than the genus g, D is winnable since the right-hand side of Riemann–Roch would be positive, so there is no way that $r(D)$ could have a negative value. For more on the exchange between this result and Theorem 2.2.5, we direct the interested reader to [BN07].

[BN07] tackles the problem of proving this theorem head on. The paper starts by defining a set $\mathcal{N} = \{D \in Div(G) \mid deg(D) = g-1, |D| = \emptyset\}$ and a map $\epsilon\colon Div(G) \to \mathbb{Z}/2\mathbb{Z}$ where $\epsilon(D) = 0$ when $|D| \neq \emptyset$ and $\epsilon(D) = 1$ when $|D| = \emptyset$. The purpose of these is to prove that the Riemann–Roch Theorem is true if and only if the following two conditions hold: (i) For every $D \in Div(G)$, there exists $v \in \mathcal{N}$ such that $\epsilon(D) + \epsilon(v - D) = 1$, (ii) For every $D \in Div(G)$ with $deg(D) = g-1$, $\epsilon(D) + \epsilon(K - D) = 0$. We will not provide a proof of Riemann–Roch on these terms, but instead offer a more streamlined proof as explicated in [Bak14].

If we put an orientation O on a graph G, we can define an associated divisor $D_O = \Sigma_{v \in V(G)}\ (indeg_O(v) - 1)(v)$, where $indeg_O(v)$ is the number of edges with v as their head in the orientation. $deg(D_O) = g - 1$ and if $\bar{O}$ is the reverse orientation of O, $D_O + D_{\bar{O}} = K_G$. We say that an orientation is *acyclic* if it contains no directed cycle. If $v = D_O$, where O is acyclic, we say v is a *moderator* and $\bar{v} = D_{\bar{O}}$ is the *dual moderator*.

Lemma 2.3.8. *If O is acyclic, D_O is not equivalent to an effective divisor.* [Bak14, Lemma 2]

Lemma 2.3.9. *Given a divisor D, either D is equivalent to an effective divisor or $v - D$ is equivalent to an effective divisor for some moderator v.* [Bak14, Lemma 3]

Lemma 2.3.10. *$r(D) = min\ deg^+(D' - v) - 1$ for every divisor D, where the minimum is taken over all divisors D' equivalent to D and all moderators v and $deg^+(D)$ is the sum of non-negative coefficients of D.* [Bak14, Corollary 1]

Proof. Let $r'(D) = min\ deg^+(D' - v) - 1$. Then if $r'(D) < r(D)$, there is an effective divisor E of degree $r'(D)$ such that $r(D - E) = -1$. Then by Lemma 2.3.9, there exists a moderator x and an effective divisor E' such that $x - D + E \sim E'$. This implies $D' - x = E - E'$ for some $D' \sim D$, so $deg^+(D' - x) \leq deg(E)$. Since $deg(E) = r'(D)$, this is a contradiction. Thus, $r'(D) \geq r(D)$. So, if we choose $D' \sim D$ and x such that we satisfy the minimum set out in our initial statement, there are effective divisors E, E' where $deg(E) = r'(D) + 1$ and $D' - x = E - E'$. This implies $D - E \sim x - E'$, which is not equivalent to any effective divisor by Lemma 2.3.8. This indicates that $r'(D) \leq r(D)$. As such, we get $r'(D) = r(D)$. $\square$

We now have enough information to formulate a proof of the Riemann–Roch Theorem for Graphs.

Proof of Theorem 2.3.19. For every moderator v, $deg^+(D - v) = deg(D) - g + 1 + deg^+(v - D)$. Let $D \sim D'$ and recall that rank is invariant over linearly equivalent divisors. By Lemma 2.3.10, we have $r(D) + 1 = min\ deg^+(D' - v) = deg(D) - g + 1 + min\ deg^+(K_G - D' - (K_G - v))$. Recall that $K_G - v = \bar{v}$, so we can continue to take minima over $D' \sim D$ and $\bar{v}$. Then $r(D) + 1 = deg(D) - g + 1 + min\ deg^+(K_G - D' - \bar{v})$. By Lemma 2.3.10 again, $min\ deg^+(K_G - D' - \bar{v}) = r(K_G - D) + 1$, so $r(D) + 1 = deg(D) - g + 1 + r(K_G - D) + 1$. Then clearly, $r(D) - r(K_G - D) = deg(D) + 1 - g$. $\square$

As a corollary to the Riemann–Roch theorem on graphs, we get a graph-theoretic analogue of another well-known result in algebraic geometry, Clifford's Theorem. We conclude this section by offering its proof. First, we define a *special* divisor D as a divisor where $|K_G - D| \neq \emptyset$.

Corollary 2.3.11. *(Clifford's Theorem for Graphs) Let D be an effective special divisor on a graph G. Then $r(D) \leq \frac{1}{2}deg(D)$.* [BN07, Corollary 3.5]

Proof. If D is an effective special divisor, then $K_G - D$ is also effective. By Lemma 2.1.7, this implies that $r(D) + r(K_G - D) \leq r(K_G)$. Note that by Theorem 2.3.19, $r(K_G) = g - 1$, so $r(D) + r(K_G - D) \leq g - 1$. Simultaneously, Theorem 2.3.19 implies that $r(D) - r(K_G - D) = deg(D) + 1 - g$. If we add these two expressions together, we get that $2r(D) \leq deg(D)$, i.e. $r(D) \leq \frac{1}{2}deg(D)$. □

2.4. **The Algebraic Turn.** Recent papers have more closely investigated the algebraic structure of the Picard group. We follow suit, as the results we prove in this thesis concern the direct-sum decomposition of $Pic_0(G_n)$, where G_n is the n^{th} cone over G. The results in this section will provide the ground work for the theorems in Section 3, as our work attempts to further sharpen the insights laid out here. Let G be a finite graph. We will start by considering a subset $S = \{w_1, ..., w_m\} \subset V(G)$. This subset has the *conformity property* when the subgraph on S is either completely disconnected or complete and for every vertex $x \notin S$, $w_i x$ is an edge in $E(G)$ if and only if $w_j x$ is an edge for all i, j. In Theorem 2.4.3, [BMZB18] uses a classic module construction $Div^0_{\mathbb{Q}}(G) = Div^0(G) \otimes_{\mathbb{Z}} \mathbb{Q}$. When $L(G)$ acts on this space, it works as a linear endomorphism. We present the statement of two Lemmas from [BMZB18] and then reproduce the proof of Theorem A from [BMZB18], generalizing to disconnected graphs. The n^{th} cone of a graph is necessarily connected, so our comments will only need to provide clarity where we are directly working with the original graph G. This will allow us to then consider the results in [GP19], which provide a different angle by which to consider the questions pertinent to this thesis.

Lemma 2.4.1. *Assume G is connected with at least 3 vertices. Suppose v_1, v_2 are a pair of vertices of degree d with the conformity property. Let $e_{12} = v_1 - v_2$ be an element of $Pic^0(G)$. Then if $v_1v_2 \in E(G)$, e_{12} has order $d+1$. Otherwise, e_{12} has order d.* [BMZB18, Lemma 3.1]

Lemma 2.4.2. *Let $j \geq 1$ and let $S^1 = \{v_1^1, \ldots, v_{m_1}^1\}, \ldots, S^j = \{v_1^j, ..., v_{m_j}^j\}$ be j mutually disjoint vertex sets, each with the conformity property. Assume G is connected and $\{S^1, ..., S^j\}$ does not completely cover G. Then if the elements $e_{1k}^i = v_1^i - v_k^i$ where $1 \leq i \leq j$ and $2 \leq k \leq m_i$ satisfy $\Sigma_{i=1}^j \Sigma_{k=2}^{m_i} \alpha_{ij} e_{1k}^i = 0$ in $Pic^0(G)$, then each $\alpha_{ij} e_{1k}^i = 0$ in $Pic^0(G)$.* [BMZB18, Lemma 3.2]

Theorem 2.4.3. *Let G be an undirected graph on $k \geq 1$ vertices. Let $n \geq 1$ be an integer, and let G_n be the n^{th} cone over G. Then there is a short exact sequence of abelian groups*

$$0 \to (\mathbb{Z}/(n+k)\mathbb{Z})^{n-1} \to Pic^0(G_n) \to H_n \to 0$$

where the order of H_n *is* $|p_{L(G)}(-n)|/n$ *where* $p_{L(G)}(x)$ *is the characteristic polynomial of* $L(G)$. [BMZB18, Theorem A]

Proof. When $k = 1$, $G_n = K_{n+1}$. It is easy to show that $Pic^0(K_{n+1})$ is isomorphic to $(\mathbb{Z}/(n-1)\mathbb{Z})^{n+1}$, so we may assume that $k \geq 2$.

Consider the matrix $B_n = L(G_n) - L(K_{n+k})$. Every entry is 0 except for the upper k by k submatrix, which we use to define $B_0 = L(G) - L(K_k)$. B_n acts on $Div^0_{\mathbb{Q}}(G_n)$, so say $\mathbf{u} \in Div^0_{\mathbb{Q}}(G_n)$ is an eigenvector of B_n with eigenvalue μ. As K_{n+k} is a complete graph, the matrix $L(K_{n+k})$ acts as multiplication by $n+k$ on all elements of $Div^0_{\mathbb{Q}}(G_n)$. Thus $\mathbf{u}$ is an eigenvector of $L(K_{n+k})$ with eigenvalue $n+k$, so it is an eigenvalue of $L(G_n)$ with eigenvalue $n+k+\mu$. This outlines a method by which we can use the eigenvectors of one of these matrices to define eigenvectors of the others.

Choose $k-1$ eigenvectors $\mathbf{u}_1, \ldots, \mathbf{u}_{k-1} \in Div^0_{\mathbb{Q}}(G)$ of B_0, with corresponding eigenvalues μ_i. Then, by appending n zeros to each vector, we get eigenvectors of $L(G_n)$ with eigenvalues $n+k+\mu_i$. Since $\mathbf{u}_1, \ldots, \mathbf{u}_{k-1}$ are eigenvectors of $L(G)$ with eigenvalues $k+\mu_i$, this is the only real place in the proof where the connectivity of G is relevant. The rank of $L(G)$ is $k-c$ where c is the number of connected components of G so when G is disconnected, there are fewer than $k-1$ nonzero eigenvalues available. In this scenario, some of the $\mathbf{u}_i$ we choose are such that $\mu_i = -k$, as their eigenvalue in $L(G)$ is zero. This does not preclude us from choosing such $\mathbf{u}_i$, but we should note that we will get some eigenvectors of $L(G_n)$ with eigenvalue n.

For $i > k+1$, the $n-1$ vectors $\mathbf{u}_i = v_i - v_{k+1}$ are eigenvectors of $L(G_n)$ with eigenvalue $n + k$ by Lemma 2.4.1. Finally, the vector $n\Sigma_{i=1}^{k} v_i - k\Sigma_{i=k+1}^{n+k} v_i$ is an eigenvector, also with eigenvalue $n+k$. This gives a basis for $Div^0_{\mathbb{Q}}(G_n)$ in eigenvectors of $L(G_n)$. By the matrix-tree theorem, the order of $Pic^0(G_n)$ is the product of these eigenvalues divided by $n+k$ since G_n is connected, which is $(n+k)^{n-1}\prod(n+k+\mu_i)$.

The elements $v_{k+1} - v_{k+1+i}$ for $i > 0$ generate a subgroup isomorphic to $(\mathbb{Z}/(n+k)\mathbb{Z})^{n-1}$ by Lemma 2.4.2. The quotient has order $\prod(n+k+\mu_i)$, but remember that the $k+\mu_i$ are eigenvalues of $L(G)$ acting on $Div^0_{\mathbb{Q}}(G)$. So, this expression is equal to $|p_{L(G)}(-n)/-n|$, where we scale by $-n$ because our $k-1$ eigenvalues correspond to all but the canonical zero eigenvalue of the Laplacian. Recall, however, that when G is disconnected, some of these $k+\mu_i$ are zero as well. $\square$

[BMZB18] poses the question: For which graphs and values of n does this short exact sequence split? In other words, when is $Pic^0(G_n)$ the direct sum of $(\mathbb{Z}/(n+k)\mathbb{Z})^{n-1}$ and H_n? [GP19] produces another direct-sum decomposition of the Picard group by directly analyzing the reduced Laplacian. They generalize to Eulerian digraphs, but for our purposes, we will assume that G is a finite and undirected (though not necessarily connected) graph as before.

Theorem 2.4.4. *Let G be a graph on $k \geq 1$ vertices with Laplacian $L(G)$. Let G_n be the n^{th} cone over G where $n \geq 2$. Let $\mathbf{1}$ be the $k \times k$ matrix with every entry equal to 1 and let I_k be the $k \times k$ identity matrix. Then $Pic^0(G_n) \cong (\mathbb{Z}/(n+k)\mathbb{Z})^{n-2} \oplus cok(nI_k + L(G) + \mathbf{1})$. Furthermore, the group $cok(nI_k + L(G) + \mathbf{1})$ has a subgroup isomorphic to $(\mathbb{Z}/(n+k)\mathbb{Z})$.* [GP19, Theorem 1]

We will not reprint the proof of this statement, mainly because it does not reference the connectivity of G, so requires no updating. The direct-sum decomposition of Theorem 2.4.4 reduces the question of when our sequence splits to an analysis of the cokernel of $nI_k + L(G) + \mathbf{1}$. Since this cokernel contains $(\mathbb{Z}/(n+k)\mathbb{Z})$ as a subgroup, we are left to wonder when this subgroup is an invariant factor. As pointed out in [GP19], this can be detected by testing the coprimality of $n+k$ and $|cok(nI_k + L(G) + \mathbf{1}|/n+k$. However, a more fruitful approach is in considering the Smith Normal Form of $nI_k + L(G) + \mathbf{1}$, as it provides the invariant factors of the cokernel in a more holistic manner. The results in the proceeding section toy with this matrix in different ways, using particular graph constructions as a means of definitively finding splitting properties. Some of our results rely on coprimality arguments, but others use graph-theoretic properties of certain G to prove conditions for splitting via information encoded in $nI_k + L(G) + \mathbf{1}$. These results allow us to extrapolate cardinality properties in Section 4 regarding how often graphs split over the n^{th} cone for given n, or over all n.

3. Proof of Results

3.1. **Theorem A.** Let G be a finite, simple, undirected graph on k vertices. We say that a vertex $v \in V(G)$ is a *universal vertex* if $vw \in E(G)$ for all $w \in V(G)$, $w \neq v$. The degree of a universal vertex is $k - 1$.

Theorem A. *If a graph G contains a universal vertex, then it splits over the n^{th} cone for all $n \geq 2$.*

Proof. Index $V(G)$ such that the universal vertex is v_1 and we have a basis to construct the Laplacian. Then

$$L(G) = \left(\begin{array}{c|c} k-1 & -\mathbf{1} \\ \hline -\mathbf{1} & \tilde{L}(G) \end{array}\right)$$

If we add nI_k and $\mathbf{1}$ to this matrix, we get

$$nI_k + L(G) + \mathbf{1} = \left(\begin{array}{c|c} n+k & \mathbf{0} \\ \hline \mathbf{0} & nI_{k-1} + \tilde{L}(G) + \mathbf{1} \end{array}\right)$$

Then clearly, $cok(nI_k + L(G) + \mathbf{1}) = \mathbb{Z}/(n+k)\mathbb{Z} \oplus cok(nI_{k-1} + \tilde{L}(G) + \mathbf{1})$ and the sequence in Theorem 2.4.3 splits. Since our choice of n was arbitrary and is only bounded below by the result from Goel, this means the sequence splits for all $n \geq 2$. $\square$

3.2. **Theorem B.** The next three sections are dedicated to examining when particular graphs split. First, we examine the 4-cycle — a graph on 4 vertices which could informally be characterized as a square.

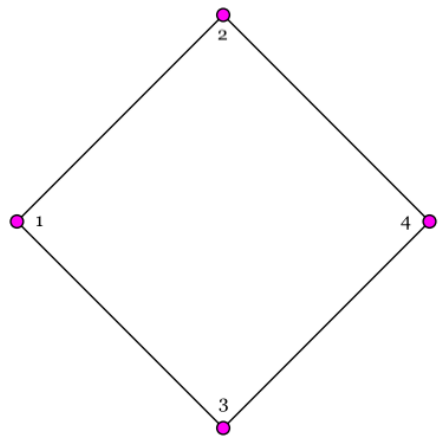

Figure 12: A 4-cycle

Theorem B. *The 4-cycle splits only over all odd cones.*

Proof. Index $V(G)$ as in the diagram above. Then

$$nI_k + L(G) + \mathbf{1} = \begin{pmatrix} n+3 & 0 & 0 & 1 \\ 0 & n+3 & 1 & 0 \\ 0 & 1 & n+3 & 0 \\ 1 & 0 & 0 & n+3 \end{pmatrix}$$

Perform the following row and column operations on $nI_k + L(G) + \mathbf{1}$: add row 4 to row 1, then subtract column 1 from column 4. This modified matrix, denoted A, has the same Smith Normal Form as $nI_k + L(G) + \mathbf{1}$, but is of the form

$$A = \begin{pmatrix} n+4 & 0 & 0 & 0 \\ 0 & n+3 & 1 & 0 \\ 0 & 1 & n+3 & 0 \\ 1 & 0 & 0 & n+2 \end{pmatrix}$$

When n is odd, $n+4 \equiv 2 \bmod n+2$. Note that $n+4$ and $n+2$ are both odd in this scenario. This means there is some $m \in \mathbb{Z}$ such that $m(n+4) \equiv -1 \bmod n+2$. Add m times row 1 to row 4. Then there is an integer multiple of column 4 that we can subtract from column 1 such that we get the matrix

$$\begin{pmatrix} n+4 & 0 & 0 & 0 \\ 0 & n+3 & 1 & 0 \\ 0 & 1 & n+3 & 0 \\ 0 & 0 & 0 & n+2 \end{pmatrix}$$

Let B denote the 3 x 3 matrix created by deleting the first row and column from A. Then $cok(nI_k + L(G) + \mathbf{1}) = \mathbb{Z}/(n+4)\mathbb{Z} \oplus cok(B)$. So, when n is odd, the sequence in Theorem 2.4.3 splits. We can show that this statement is "if and only if" by directly computing the Smith Normal Form of $cok(nI_k + L(G) + \mathbf{1})$.

Since the Smith Normal Form of a matrix is a diagonal matrix, we can denote its entries by the set $\{d_1, d_2, d_3, d_4\}$. Each d_i is the greatest common denominator of all i x i minors of $nI_k + L(G) + \mathbf{1}$, divided by d_{i-1}. Additionally, $d_i | d_{i+1}$. Since the top right 2 x 2 minor of $nI_k + L(G) + \mathbf{1}$ is -1, d_1 and d_2 both are 1. It is easy to calculate $d_3 = (n+2)(n+4)$ from the 3 x 3 minors of $nI_k + L(G) + \mathbf{1}$. Then d_4 is $(n+2)(n+4)$ as well. We use a coprimality argument here — $n+4$ never factors out from the cokernel when n is even, as $n+2$ and $n+4$ share a factor of 2 (this is an application of the Fundamental Theorem of Finite Abelian Groups). When n is odd, $n+4$ always factors out, as it is always coprime with $n+2$. Thus, the sequence in Theorem 2.4.3 splits on the 4-cycle if and only if n is odd. □

3.3. **Theorem C.**

Theorem C. *The path graph on 4 vertices splits only over all odd cones.*

Proof. Put a "left-to-right" indexing on $V(G)$. Then

$$nI_k + L(G) + \mathbf{1} = \begin{pmatrix} n+2 & 0 & 1 & 1 \\ 0 & n+3 & 0 & 1 \\ 1 & 0 & n+3 & 0 \\ 1 & 1 & 0 & n+2 \end{pmatrix}$$

Perform the following row and column operations on $nI_4 + L(G) + \mathbf{1}$, add row 1 to row 4, subtract column 4 from column 1, add row 2 to row 3, subtract column 3 from column 2, add column 2 to column 3, subtract row 3 from row 2, add row 4 to row 3, subtract column 3 from column 4, add column 2 to column 4 We get a matrix A

with the same Smith Normal Form as $nI_4 + L(G) + \mathbf{1}$ that looks like

$$\begin{pmatrix} n+1 & -1 & 0 & 0 \\ -1 & n+3 & 0 & n+3 \\ 0 & 0 & n+4 & 0 \\ 0 & 0 & 1 & n+2 \end{pmatrix}$$

We can use a similar method as we did in Section 3.2 to reduce the bottom right 2×2 matrix to

$$\begin{pmatrix} n+1 & -1 & 0 & 0 \\ -1 & n+3 & 0 & n+3 \\ 0 & 0 & n+4 & 0 \\ 0 & 0 & 0 & n+2 \end{pmatrix}$$

when n is odd. Then, subtract row 4 from row 2, add column 1 to column 4, subtract row 1 from row 4, add $n+1$ times column 4 to column 1, subtract column 4 from column 2, and subtract $n+1$ times row 4 from row 1. This gives us the matrix

$$B = \begin{pmatrix} (n+2)(n+1) & -(n+2) & 0 & 0 \\ -1 & n+3 & 0 & 0 \\ 0 & 0 & n+4 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

Let C denote the upper left 2 x 2 submatrix. Then $\text{cok}(nI_k + L(G) + \mathbf{1}) = \mathbb{Z}/(n+4)\mathbb{Z} \oplus cok(C)$. So, when n is odd, we can use row and column operations to show that the sequence in Theorem 2.4.3 splits. During the production of this thesis, [GP19] published an updated version of their paper, which now contains an algebraic proof of the “only if” direction of this result. They utilize a similar approach as we do in Theorem B. For this reason, we will not analyze the Smith Normal Form of $nI_4 + L(G) + \mathbf{1}$ here, but instead refer the interested reader to their work. $\square$

Figure 13: The path graph on 4 vertices

3.4. **Theorem D.** If we treat n as a variable over which we analyze the splitting properties of $Pic^0(G_n)$, we can think of $nI_k + L(G) + \mathbf{1}$ as a matrix over the polynomial ring of $\mathbb{Q}$. The Smith Normal Form of a matrix is commonly presented as $SNF(M) = PMA$, where P is the product of elementary row matrices and A is the product of elementary column matrices. Smith normalization is derived through a series of row and column operations which first diagonalizes a matrix and then imposes our divisibility conditions upon it so these matrices P, A keep track of those manipulations. When P and A both have coefficients entirely in $\mathbb{Z}$, they provide an accurate account of the normalization process with respect to our Picard group. This is because we treat the Laplacian as an operator over $\mathbb{Z}^n$, so Pic^0 really keeps track of integer relations between degree zero divisors. Allowing for rational row/column operations creates distortions, such as rescaling that alters the order of $cok(nI_k + L(G) + \mathbf{1})$.

Working over $\mathbb{Q}[x]$ allows for fast algorithms in computing the Smith Normal Form, but rarely do we find graphs such that $SNF(nI_k + L(G) + \mathbf{1})$ has associated integer matrices P and A, rather than rational matrices. This proof considers a case in which P and A are integral.

Theorem D. *The 5-cycle splits only over the nth cone for $n \not\equiv 0$ mod 5.*

Proof. When G is the 5-cycle as depicted in the diagram below, the normalization process for $nI_5 + L(G) + \mathbf{1}$ is as follows:

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & n^2+5n+5 & 0 \\ 0 & 0 & 0 & 0 & n^3+10x^2+30n+25 \end{pmatrix} =$$

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ -1 & -n-2 & n+2 & 1 & 0 \\ n^2+6n+6 & -n^2-7n-9 & n^2+6n+6 & 1 & 1 \end{pmatrix} \begin{pmatrix} n+3 & 0 & 1 & 1 & 0 \\ 0 & n+3 & 0 & 1 & 1 \\ 1 & 0 & n+3 & 0 & 1 \\ 1 & 1 & 0 & n+3 & 0 \\ 0 & 1 & 1 & 0 & n+3 \end{pmatrix} \begin{pmatrix} 0 & 0 & 0 & -1 & 1 \\ 1 & -1 & 1 & 0 & -n-4 \\ 1 & -1 & 1 & 1 & -n-4 \\ 0 & 1 & -1 & n+2 & 1 \\ -n-3 & n+3 & -n-2 & -n-2 & n^2+7n+11 \end{pmatrix}$$

Observe $d_5 = n^3+10n^2+30n+25$ in the Smith Normal Form. This is a polynomial in n that can be factored as $(n^2+5n+5)(n+5)$. Our splitting question, then, is reduced to analyzing for which n these two factors are coprime. Since $(n^2+5n+5) = n(n+5)+5$, this can be further reduced to question of when $n+5$ and 5 are coprime. Clearly, this occurs if and only if $gcd(n,5) = 1$, so our sequence splits for the 5-cycle if and only if $n \not\equiv 0$ mod 5. $\square$

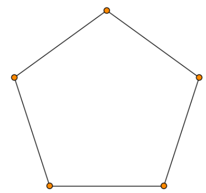

Figure 14: A 5-cycle

3.5. Theorem E.

Theorem E. *The totally disconnected graph on k vertices splits only over the n^{th} cone when $gcd(n,k) = 1$.*

Proof. Index $V(G)$ arbitrarily. The Laplacian of the totally disconnected graph on k vertices is just the k x k zero matrix. As such, $nI_k + L(G) + \mathbf{1}$ is just the k x k matrix with $(n+1)$'s on the diagonal and 1's everywhere else. Perform the following row operations on $nI_k + L(G) + \mathbf{1}$ simultaneously:

1. Subtract the first row from all rows but the first and last.
2. Add every row but the first and last to the last row.
3. Subtract $(x + k - 1)$ copies of the first row from the last row.

These row operations on $nI_k + L(G) + \mathbf{1}$ produce the matrix:

$$P' = \left(\begin{array}{c|c|c} n+1 & \mathbf{1} & 1 \\ \hline -\mathbf{x} & nI_{k-2} & \mathbf{0} \\ \hline -x(x+k) & \mathbf{0} & 0 \end{array}\right)$$

In the context of the totally disconnected graph, we have effectively produced a "generalized" P where $SNF(nI_k + L(G) + \mathbf{1}) = P(nI_k + L(G) + \mathbf{1})A$. We pair this with a "generalized" A, which is given by the following algorithm. Perform these column operations on P' simultaneously, as before:

1. Multiply the first column by zero and the last column by $(x + k + 1)$.
2. Add the last column to the first column.
3. Subtract the last column from each but the first and last column.
4. Subtract all but the last column from the last column.

These column operations on P' produce the matrix:

$$\left(\begin{array}{c|c|c} 1 & \mathbf{0} & 0 \\ \hline \mathbf{0} & nI_{k-2} & \mathbf{0} \\ \hline 0 & \mathbf{0} & n(n+k) \end{array}\right)$$

This matrix is the Smith Normal Form of $nI_k + L(G) + \mathbf{1}$. Clearly, the sequence in Theorem 2.4.3 splits for the totally disconnected graph if and only if $gcd(n, k) = 1$. $\square$

3.6. **Theorem F.**

Theorem F. *Given a finite graph G, there are infinitely many $n \in \mathbb{Z}$ such that G splits over the n^{th} cone.*

Proof. As a direct consequence of Theorem 2.4.3, we get that $|Pic^0(G_n)| = (n+k)^{n-1}|p_{L(G)}(-n)/-n|$ (See Corollary B, [BMZB18]). As we have indicated throughout this section, one way to detect that the sequence in Theorem 2.4.3 splits is through the coprimality of $(n+k)$ and $|cok(nI_k + L(G) + \mathbf{1})|/(n+k)$. Another would be to test the coprimality of $(n+k)$ and $|p_{L(G)}(-n)/-n|$. Since these values correspond to the order of subgroups of $Pic^0(G_n)$ and their product is the order of $Pic^0(G_n)$, their coprimality would indicate that $Pic^0(G_n) = (\mathbb{Z}/(n+k)\mathbb{Z})^{n-1} \oplus H_n$. Thus, we can assess whether on an arbitrarily given G, there exist values of n such that these two orders are coprime.

Let $n = p - k$ for some prime $p > |p_{L(G)}(k)|$. Then $(n+k) = p$. Consider $|p_{L(G)}(k-p)/k-p|$. If we factor out $p_{L(G)}(k-p)$ such that we can separate the terms with a factor of $(-p)$ from the rest of the polynomial, we are left with $p_{L(G)}(k)$ plus something that is a multiple of p. Since $p > |p_{L(G)}(k)|$ by assumption, this implies that $p_{L(G)}(k-p)$ is not a multiple of p. Then neither is $|p_{L(G)}(k-p)/k-p|$, so $(n+k) = p$ and $|p_{L(G)}(k-p)/k-p|$ are coprime. This indicates that for $p > |p_{L(G)}(k)|$, the sequence in Theorem 2.4.3 splits when $n = p-k$. Since G is finite by assumption, $|p_{L(G)}(k)|$ is finite as well, so there are infinitely many p such that this condition holds. $\square$

4. Applications and Further Speculation

This thesis considers the direct-sum decompositions of the Picard group of a graph. When the sequence in Theorem 2.4.3 does not split, it implies that the subgroup $(\mathbb{Z}/(n+k)\mathbb{Z}) \subset cok(nI_k + L(G) + \mathbf{1})$ is embedded in a larger cyclic subgroup of $cok(nI_k + L(G) + \mathbf{1})$. Returning to our chip-firing intuitions, this says that some divisor with 1 on one cone vertex, -1 on another, and zero elsewhere is linearly equivalent to an integer multiple of some other divisor in $Pic^0(G_n)$. Our concern, then, is with the situations in which the divisors $\{v_{k+1} - v_{k+1+i}\}$ for $0 < i < n$ are independent from the rest of the elements in $Pic^0(G_n)$, i.e. H_n. This can happen in two ways. When we take the Smith Normal Form of $cok(nI_k + L(G) + \mathbf{1})$, we may find that $(n + k)$ is one of the invariant factors on the diagonal. Alternatively, as was commonly seen in our results, $(n + k)$ might be a factor of some larger invariant factor. One avenue that we didn't have the time to pursue further was formalizing this difference. If we consider the former a "strong split" ($n + k$ is an invariant factor regardless of n) and the latter a "weak split" ($n+k$ is an invariant factor for good values of n), there are a few questions we can ask. Does the existence of one versus the other imply properties about the graph under consideration? Can we locate graph properties that produce one versus the other? This second question was partially answered by our Theorem A, which showed that we get a strong split as such when G has a universal vertex. Hopefully, such a conceptual framework can create a more subtle understanding of chip-firing games.

Our results from the last section can also be used to prove how many graphs split over the n^{th} cone for a given n. Our result in Section 3.1 proves two things. First, it indicates that there are infinitely many graphs that split over every cone. This is because there are infinitely many graphs that have a universal vertex. This dually implies that given some value n, there are infinitely many graphs that split over the n^{th} cone. On the other hand, there are also infinitely many graphs which exist that do not split over the n^{th} cone. This is implied by Section 3.5. Given n, there are infinitely many k such that $gcd(n, k) \neq 1$, each corresponding to a totally disconnected graph on k vertices. By our result, these graphs all do not split over the n^{th} cone. If we were to ask if infinitely many graphs never split, this would be easily contradicted by Section 3.6, which not only states that every graph splits at some point, but that every graph does so infinitely many times. It would be interesting to know if every graph also doesn't split infinitely many times, or if there is a family of graphs which doesn't split only finitely often. Since our universal vertex graphs satisfy this trivially, we can additionally require that our sequence doesn't split at least once.

We often relied on a coprimality argument to indicate when graphs would carry a splitting property. Indeed, one of the insights of this thesis (hopefully) was in treating the matrix $nI_k + L(G) + \mathbf{1}$ as a matrix over the polynomial ring. Once we had normalized this polynomial matrix, we could produce results through simple coprimality observations. However, Section 3.1 broke from this tradition and introduced a graph-theoretic perspective on the Picard group. The hope is that our Section 3.1 can produce other graph-theoretic approaches to understanding the structure of Pic^0. $nI_k + L(G) + \mathbf{1}$ is not structurally very different from the graph Laplacian and as such, it might encode more fruitful information about our graph of choice. One other consideration that wasn't pursued further regarded the relationship between connectivity and splitting. In practice, it appeared that the more connected a graph G was, the more likely it was to split more often over iterated cones. This sort of insight can be gleaned from Section 3.1, as when G becomes more connected, diagonal entries of $nI_k + L(G) + \mathbf{1}$ tend towards $n + k$ and nondiagonal entries tend more frequently to zero. Once again, we were not able to formalize this by the time of writing. Certain sources, namely [Mer99], have related critical configurations in the chip-firing game in [Big99] to the Tutte polynomial, and perhaps this is a graph invariant that can relate the Picard group and graph connectivity more directly. All in all, we offer a modest contribution to the study of chip-firing games, one that can hopefully sharpen already-present techniques while drawing new connections. An object which powerfully draws together such disparate fields of mathematics under the guise of a compelling, solitaire game is worthy of much further attention.

References

[Bak14] Matthew Baker, *Reduced divisors and Riemann–Roch for Graphs*, 2014. ↑38, 39

[Bak] ———, *Reduced Divisors on Graphs*. ↑16

[Bak08] ———, *Specialization of linear systems from curves to graphs*, Algebra & Number Theory **2** (2008), no. 6, 613–653. With an appendix by Brian Conrad. MR2448666 (2010a:14012) ↑31, 32, 33, 34, 35, 36

[Bak10] ———, *The Dollar Game on a Graph (With 9 Surprises!)*, 2010. ↑

[BN07] Matthew Baker and Serguei Norine, *Riemann–Roch and Abel–Jacobi theory on a finite graph*, Adv. Math. **215** (2007), no. 2, 766–788. ↑7, 15, 16, 18, 28, 36, 37, 38, 40

[Big97] Norman L. Biggs, *Algebraic potential theory on graphs*, Bull. London Math. Soc. **29** (1997), 641–682. ↑29

[Big99] ———, *Chip-Firing and the Critical Group of a Graph*, Journal of Algebraic Combinatorics **9** (1999), no. 1, 25-45. ↑10, 18, 23, 25, 26, 28, 29, 30, 57

[BLS91] Anders Björner, László Lovász, and Peter W. Shor, *Chip-firing games on graphs*, European J. Combin. **12** (1991), no. 4, 283–291. ↑18, 19, 20, 21, 22, 23

[Bol98] Bela Bollobas, *Modern Graph Theory*, Graduate texts in mathematics, Springer, 1998. ↑

[BMZB18] Morgan Brown, Jackson S. Morrow, and David Zureick-Brown, *Chip-Firing Groups of Iterated Cones*, Linear Algebra & its Applications **556** (2018), 46–54. ↑4, 5, 6, 30, 40, 41, 42, 43, 54

[DF04] David S. Dummit and Richard M. Foote, *Abstract Algebra*, Wiley, 2004. ↑11, 31

[GP19] Gopal Goel and David Perkinson, *Critical groups of iterated cones*, Linear Algebra & its Applications **567** (2019), 138–142. ↑5, 6, 30, 41, 43, 44, 49

[Har18] Kevin Hartnett, *Tinkertoy Models Produce New Geometric Insights*, 2018. ↑30

[Har77] Robin Hartshorne, *Algebraic Geometry*, Graduate Texts in Mathematics, Springer, 1977. ↑31

[Lor08] Dino Lorenzini, *Smith normal form and Laplacians*, J. Combin. Theory Ser. B **98** (2008), no. 6, 1271–1300. ↑

[Mer99] Criel Merino, *Matroids, the Tutte Polynomial, and the Chip Firing Game*, Somerville College, University of Oxford, 1999. ↑57

[Mer98] Russell Merris, *Laplacian graph eigenvectors*, Linear Algebra & its Applications **278** (1998), no. 1, 221–236. ↑

[Nel] Peter Nelson, *An Introduction to Schemes*. ↑31

[Sta13] Richard P. Stanley, *Algebraic combinatorics*, Undergraduate Texts in Mathematics, Springer, New York, 2013. Walks, trees, tableaux, and more. ↑11

[Vak] Ravi Vakil, *Introduction to Algebraic Geometry, Class 5*. ↑31