

GEOMETRIC BIJECTIONS OF GRAPHS AND REGULAR MATROIDS

A Dissertation
Presented to
The Academic Faculty

By

Chi Ho Yuen

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in
Algorithms, Combinatorics, and Optimization

School of Mathematics
Georgia Institute of Technology
May 2018

Copyright © Chi Ho Yuen 2018

GEOMETRIC BIJECTIONS OF GRAPHS AND REGULAR MATROIDS

Approved by:

Professor Matthew Baker, Advisor
School of Mathematics
Georgia Institute of Technology

Professor Josephine Yu
School of Mathematics
Georgia Institute of Technology

Professor Prasad Tetali
School of Mathematics
Georgia Institute of Technology

Professor Xingxing Yu
School of Mathematics
Georgia Institute of Technology

Professor Richard Peng
School of Computer Science
Georgia Institute of Technology

Date Approved: , 2018

To my family.

To Diocesan Boys' School, as I have promised more than ten years ago.

ACKNOWLEDGEMENTS

First of all, I am grateful to my advisor Matthew Baker for providing me a tremendous amount of support and advice throughout my graduate years. He has introduced me to many interesting subjects and problems in mathematics, and mentored my exploration of these topics. Besides mathematics itself, he has also taught me a lot in other essential aspects of being a mathematician.

Then I would like to thank Robin Thomas for giving me a chance to join the amazing ACO program; I would say that this dissertation utilizes material I have learnt from almost all of the seven core courses together with several other ACO courses. I also thank other faculty from the ACO program and the School of Mathematics for their guidance, including but certainly not limited to Josephine Yu, Xingxing Yu, Prasad Tetali, Richard Peng, Santosh Vempala, Tom Trotter, and Cathy Jacobson.

Moreover, I am very lucky to have met many supportive friends in Georgia Tech, who helped me to get through . The long list include Marcel Celaya, Haiyu Zhou, Shane Scott, Elizabeth Bolduc, Dario Mena, Hagop Tossounian, Sudipta Kolay, Yan Wang, Chun Hung Liu, Guo He, Juntao Duan, Dawei He, Ruidong Wang, Thanh Dang, Benjamin Cousins, James Bailey, Samantha Petti, Rose McCarty, Charles Wang, and Grace Harper. I also thank the people of Georgia Tech Chess Club, especially Jason Wu for giving me a chance to play for Georgia Tech in the Georgia Collegiate Team Championship.

Outside Georgia Tech, I need to thank Spencer Backman, Farbod Shokrieh, David Perkinson, Caroline Klivans, Melody Chan, Olivier Bernardi, Greg Kuperberg, Lionel Levin, Victor Reiner, Francisco Santos, Emeric Gioan, Jan Draisma, Sam Hopkins, Bo Lin, Swee Hong Chan, Ralph Morrison, Yoav Len, and Alexander Lazar for insightful discussions.

It has been fortunate in my life that even when I am in the US, I am still able to meet,

on academic occasions, many old friends from Hong Kong; their experience and encouragement motivated me to hang on from so far away from home. They include Jasper Lee, Kenneth Hung, Sze Yiu Chau, Yin Tat Lee, Mike Wong, Peter Lam, Wai Kit Lam, Ivan Li, Ken Wong, Wayne Lam, Harry Cheung, Victor Fong, Tony Wong, Marco Cheung, Pak Hin Li, and Henry Tsang.

TABLE OF CONTENTS

Acknowledgments	v
List of Figures	x
Chapter 1: Introduction	1
Chapter 2: Preliminaries	6
2.1 Graphs	6
2.2 Divisors and Chip-firing	7
2.3 Matroids and Oriented Matroids	10
2.3.1 Matroids	10
2.3.2 Oriented Matroids	13
2.3.3 Regular Matroids	15
2.3.4 Circuit and Cocircuit Lattices	18
2.4 Jacobian Groups	19
2.5 Tutte Polynomials	20
Chapter 3: Combinatorics of Break Divisors and the ABKS Decomposition . . .	22
3.1 Definitions	22
3.1.1 Metric Graphs	22

3.1.2	Break Divisors	23
3.1.3	Tropical Jacobians and the ABKS Decomposition	24
3.2	Checking Break Divisors by Matroid Intersection	28
3.3	Some Geometry of $\text{Div}_+^g(\Gamma) \rightarrow \text{Jac}(\Gamma)$	29
3.4	The f -Vector of the ABKS Decomposition	31
3.5	A Recursive Proof of the Number of (Generalized) Break Divisors	34
Chapter 4: Circuit-cocircuit Reversal Systems		38
4.1	Definition	38
4.2	Circuit-cocircuit Reversal System as a $\text{Jac}(M)$ -Torsor	39
4.2.1	Description of the Action	40
4.2.2	The Action is Well-defined	41
4.2.3	The Action is Simply Transitive	43
4.2.4	The Action is Efficiently Computable	45
4.3	Circuit-cocircuit Reversal Systems of Non-Regular Matroids	47
4.3.1	Reduction via Circuit-cocircuit Minimal Orientations	47
4.3.2	Proof of the Inequality	51
Chapter 5: Geometric Bijections		55
5.1	Statement of the Main Theorem	55
5.2	Zonotopes	57
5.3	Zonotopal Tilings	60
5.4	Proof of the Main Theorem	69
5.4.1	Proof for Realizable Oriented Matroids	69

5.4.2	Specialization to Regular Matroids	72
5.5	Ehrhart Theory	74
5.6	Group Action–Tiling Duality	79
5.7	The ABKS Decomposition as a Zonotopal Tiling	82
5.8	Algorithmic Aspects	85
5.8.1	Inverse Algorithm for Geometric Bijections via Linear Programming	85
5.8.2	Sampling Algorithms	86
5.8.3	Edge Ordering Maps and Their Combinatorial Inverses	89
5.9	A Note on Lawrence Polytopes and Lawrence Ideals	92
Chapter 6: Bernardi Processes		97
6.1	Setup of Bernardi Processes	97
6.2	Bernardi Bijections as Geometric Bijections	99
6.2.1	Planar Case	99
6.2.2	Non-planar Case	106
6.3	Bernardi Bijections as Burning Algorithms	110
6.3.1	Cori–Le Borgne Processes	110
6.3.2	Statement and Proof	111
Chapter 7: Extensions		114
7.1	Geometric Bijections for General Oriented Matroids	114
7.2	Simplicial and Cellular Trees	121
References		135

LIST OF FIGURES

2.1	Example of chip-firing moves.	8
2.2	Fano matroid.	12
3.1	Examples of chip-firing on a metric graph.	23
3.2	A fixed model for a metric graph Γ and a break divisor $(p_1) + (p_2)$ on it. . .	26
3.3	The ABKS decomposition of $\text{Pic}^2(\Gamma)$ ([1, Figure 1]), using the model chosen in Figure 3.2.	27
4.1	Example of the torsor. Here the reference orientations of e, f are the same as the orientation we begin with.	40
4.2	Illustration of Proposition 4.2.6.	43
5.1	Diagrams for Example 5.1.6. The top row specifies a (non-acyclic) cycle signature. The bottom row shows β_{σ, σ^*} need not to be bijective when the signatures failed to be acyclic.	58
5.2	The zonotope associated to M_{K_3} . We note that it is living in \mathbb{R}^3 as M_{K_3} is not of full row rank, and we use the line segment from \mathbf{e}_i to \mathbf{e}_j instead of from $\mathbf{0}$ to $\mathbf{e}_j - \mathbf{e}_i$, but none of these differences in convention changes the geometry of zonotope.	59
5.3	The zonotope associated to K_3 , with points identified as continuous circuit reversal classes. The classes corresponding to the vertices of the zonotope are acyclic, hence they are singletons, whereas the lattice point in the middle contains two (discrete) orientations as well as a family of continuous orientations between them.	61

5.4	The zonotope associated to K_3 with points identified as continuous orientation classes as in Figure 5.3. However, for the lattice point in the middle, we use the σ induced by the total ordering and reference orientation on the right to pick out a σ -compatible (continuous) orientation, which is highlighted in blue.	63
5.5	More examples of σ -compatible continuous orientations of the zonotope of K_3 . The bi-oriented elements are highlighted in red.	64
5.6	The subdivision of the zonotope associated to K_3 as described in Theorem 5.3.6 using σ induced by the total ordering and reference orientation on the right. The red edges are bi-oriented.	67
5.7	Geometric picture of the bijection for K_3 using the pair (σ, σ^*) induced by the total ordering and reference orientation from Figure 5.6.	70
6.1	Bernardi processes on two different spanning trees, using the same (planar) ribbon structure and starting pair ([13, Figure 2]).	98
6.2	Illustration of Proposition 6.2.1. Picture by Emma Cohen.	100
6.3	Some conventions on the orientation of plane graphs. Here the orientation of e induces the orientation of the dual edge e^* , and F is said to be the face on the right of (v, e)	101
6.4	BNG I (left) and BNG II (right).	107

SUMMARY

The Jacobian of a graph, also known as the sandpile group or the critical group, is a finite abelian group associated to the graph; it has been independently discovered and studied by researchers from various areas. By the Matrix-Tree Theorem, the cardinality of the Jacobian is equal to the number of spanning trees of a graph. In this dissertation, we study several topics centered on a new family of bijections, named the geometric bijections, between the Jacobian and the set of spanning trees. An important feature of geometric bijections is that they are closely related to polyhedral geometry and the theory of oriented matroids despite their combinatorial description; in particular, they can be generalized to Jacobians of regular matroids, in which many previous works on Jacobians failed to generalize due to the lack of the notion of vertices.

In Chapter 3, we study the combinatorics of break divisors and the ABKS decomposition in tropical geometry. In Chapter 4, we construct a canonical simply transitive group action of the Jacobian on the circuit-cocircuit reversal system of a regular matroid; we also study the case of non-regular oriented matroids. In Chapter 5, which is joint work with Spencer Backman and Matthew Baker, we introduce geometric bijections between the circuit-cocircuit reversal system and the set of bases of a regular matroid, and prove their bijectivity using constructions involving zonotopes and their tilings; we then further develop the theory from different perspectives, including Ehrhart theory and algorithms. In Chapter 6, we study Bernardi process on embedded graphs, and provide a connection between planar Bernardi processes and the ABKS decomposition via geometric bijections. In Chapter 7, we mention a few extensions of our work.

CHAPTER 1

INTRODUCTION

Jacobians and zonotopes are two classical and prominent objects in algebraic and geometric combinatorics, with numerous connections with other parts of mathematics. The central subject of this dissertation is the notion of *geometric bijections*, which relates the two topics in a novel manner. In the process of developing the notion, we build connections with other subjects, and as a result, we unify some concepts and results from various origins into the same framework. A recurring paradigm here is that even if one is only interested in graphs or problems that are defined combinatorially, it is often more appropriate to consider the problems in a matroidal or (polyhedral) geometric setting.

The *Jacobian group* (or Jacobian for short) of a graph, also known as the sandpile group or the critical group, is a well-known and interesting finite group invariant. It has been discovered independently by researchers in statistical physics [43], number theory [83], and combinatorics [21]. The Jacobian is closely related to the *abelian sandpile model* in physics [8] and *chip-firing game* in combinatorics [24]. Jacobians, sandpiles and their generalizations have been studied from different perspectives, such as algebraic and tropical geometry [11, 33, 53], algebra [16, 62, 93], combinatorics [12, 65, 86], and probability theory [70, 112], just to name a few.

A *zonotope* is a Minkowski sum of line segments. Three-dimensional zonotopes, also known as zonohedra, were first introduced back in the nineteenth century [50]. High dimensional generalization of zonohedra were considered by discrete geometers [37, 38, 85], and later appeared in algebra [66], algebraic geometry [2], computer science [18], etc. A major significance of zonotopes is that they capture the combinatorics of (realizable) *oriented matroids* [23, 115], an abstraction of point configurations and hyperplane arrangements in real Euclidean space. In particular, the *duality* of oriented matroids, which generalizes plane

duality in graph theory and linear programming duality in optimization and polyhedral combinatorics, is a common theme behind many results studied in this dissertation.

By a variant of Kirchhoff’s Matrix–Tree Theorem [76], the cardinality of the Jacobian is equal to the number of spanning trees of a graph, and the problem of finding explicit combinatorial bijections between the Jacobian and the set of spanning trees of a graph has received a lot of attention in combinatorics [20, 32, 34, 55, 94]. On the other hand, Jacobians can be generalized and be defined for *regular matroids* [87], with an analogous generalization of the Matrix-Tree Theorem stating that the cardinality of the Jacobian is equal to the number of *bases* of a regular matroid. However, most of the aforementioned bijections no longer generalize because they use the notion of vertices of graphs in an essential manner. In contrast, we built on the work of Gioan [57, 58] and Backman [4, 5] relating elements of Jacobian and orientation classes of a graph (resp. regular matroid), and thus we were able to formulate a family of bijective maps that is well-defined for regular matroids. Moreover, while the maps are very simple to define, proving that they are indeed bijections is non-trivial, and all known proofs use polyhedral geometry or closely related techniques, hence the name “geometric bijections”. In particular, zonotopes and their tilings play an important role in the proof.

For the purpose of referencing within this chapter, we state here the central theorem of this dissertation, together with a brief explanation of the essential notations. A more detailed discussion will be provided in later chapters.

Theorem 1.0.1 *Let M be a regular matroid and let $\mathcal{B}(M)$ be the set of bases of M . The circuit-cocircuit reversal system $\mathcal{G}(M)$ of M is the set of equivalence classes of orientations of M with respect to the equivalence relation generated by circuit and cocircuit reversals. Fix a pair of acyclic signatures (σ, σ^*) induced by a generic vector \mathbf{w} , which, for each circuit (resp. cocircuit) C of M , picks the orientation of C such that $\mathbf{w} \cdot C > 0$.*

Given a basis $B \in \mathcal{B}(M)$, let $\mathcal{O}(B)$ be the orientation of M in which we orient each $e \notin$

B according to its orientation in $\sigma(C(B, e))$ and each $e \in B$ according to its orientation in $\sigma^*(C^*(B, e))$, where $C(B, e)$ (resp. $C^*(B, e)$) is the fundamental circuit (resp. cocircuit) of e with respect to B . Then the map $B \mapsto [\mathcal{O}(B)]$ gives a bijection $\beta_{\sigma, \sigma^*} : \mathcal{B}(M) \rightarrow \mathcal{G}(M)$.

The high-level strategy of our approach is as follows. We will first describe a (combinatorial) simply transitive group action of the Jacobian $\text{Jac}(M)$ on $\mathcal{G}(M)$ in Section 4.2. This reduces the problem of finding bijections between $\text{Jac}(M)$ and $\mathcal{B}(M)$ to finding bijections between $\mathcal{G}(M)$ and $\mathcal{B}(M)$. Then we will prove the bijectivity of the maps β_{σ, σ^*} in Chapter 5 by giving them a geometric interpretation: we consider a *zonotope* Z_M associated to M and construct a *tiling* of Z_M induced by σ , which partitions Z_M into cells that canonically correspond to the bases of M . Next we *shift* the lattice points of Z_M in a direction induced by σ^* , so that each cell contains a unique shifted lattice point in its interior, we then have a bijection between cells (as bases) and a subcollection of lattice points (which can be canonically identified with the elements of $\mathcal{G}(M)$). Finally, we will show that such bijection coincides with β_{σ, σ^*} .

We highlight a few interesting ideas from our picture. First of all, while the description of β_{σ, σ^*} 's is symmetric with respect to circuits and cocircuits, the tiling versus shifting setup in the geometric picture is not; this is an instance of the oriented matroid duality, together with the realizable case of the Bohne–Dress theorem [26, 96]. Secondly, the definition of acyclic signatures generalizes the data of edge ordering (Example 5.1.4), and many results previously formulated using edge ordering can be extended to such setting, this develops the idea from the work of Greene and Zaslavsky [59]. In fact, we will prove new result (Theorem 5.5.7) that was formulated using edge ordering, but its proof relies on the general geometric setting. Finally, we give a unifying interpretation of several works on bijective problems concerning Jacobians: in the sense of *Group Action–Tiling Duality* (Theorem 5.6.3), the isomorphism of group actions induced by *Bernardi bijections* (resp. *rotor-routings*) on a plane graph G [13, 30, 31] is dual to the *ABKS decomposition* of the planar dual of G [1], cf. Theorem 6.2.4. This provides a rather surprising connection be-

tween two “canonical” constructions (cf. the question by Jordan Ellenberg during the AIM workshop on chip-firing [67]) motivated by bijective combinatorics and tropical geometry, respectively.

Now we outline the remaining chapters of this dissertation.

In Chapter 2, we review definitions and basic properties of important objects that run through the whole dissertation, including graphs and their divisors, matroids and oriented matroids, regular matroids and their circuit/cocircuit lattices, Jacobian groups, and the Tutte polynomial.

In Chapter 3, we review the notion of *break divisors* on metric graphs, followed by the construction of the ABKS decomposition of *tropical Jacobians*. Then we revisit the definition of break divisors from both combinatorial and geometric points of view, and explain how they are related in an intriguing way via tropical geometry, which might serve as an example of a potentially more general theory. Lastly, we prove two enumerative results concerning these objects, with motivations and applications in algebraic geometry. We note that trying to understand the combinatorics of these objects leads to the main work of this dissertation.

In Chapter 4, we study Gioan’s circuit-cocircuit reversal system $\mathcal{G}(M)$ and prove that $\mathcal{G}(M)$ is canonically a $\text{Jac}(M)$ -torsor when M is a regular matroid. We also study $\mathcal{G}(M)$ when M is non-regular, and show that the equality $|\mathcal{G}(M)| = |\mathcal{B}(M)|$ actually characterizes regular matroids; in the proof, we use the notion of *circuit-cocircuit minimal orientations*, which will be further generalized and play a crucial role in the rest of the dissertation.

In Chapter 5, we start with constructing zonotopes and their tilings using an approach inspired by Chapter 3, and use them to prove Theorem 1.0.1. Then we further develop the theory from different aspects: in Section 5.5, we elaborate the combinatorial interpretation of lattice points in zonotopes and give a finer version of Stanley’s formula concerning the *Ehrhart polynomial* of a unimodular zonotope; in Section 5.6, we formulate and prove a

Group Action–Tiling Duality connecting two seemingly orthogonal phenomena, namely, isomorphism of $\text{Jac}(M)$ -actions and translation invariance of tilings of unimodular zonotopes; in Section 5.7, we relate the ABKS decomposition with zonotopal tilings; in Section 5.8, we present several algorithmic consequences from the geometric bijection picture; finally, in Section 5.9, we review some *Lawrence type constructions* in polyhedral geometry and commutative algebra, and explain how our theory is related to them. Most of this chapter is joint work with Spencer Backman and Matthew Baker.

In Chapter 6, we investigate the combinatorial bijections introduced by Bernardi on embedded graphs. We prove a conjecture by Baker that Bernardi bijections (of the first type) on plane graphs are geometric bijections, and demonstrate how this connection clarifies and simplifies several previous known results on Bernardi torsors; we also prove a partial converse in the non-planar situation. Then we look at Bernardi bijections of the second type, which are more in line with the classical *burning algorithms* and do not have a geometric structure, but curiously some proof techniques from geometric bijections still apply.

In the first half of Chapter 7, we offer a generalization of Theorem 1.0.1 using more abstract oriented matroid theory, notably *oriented matroid programming*. It is part of the on-going work joint with Spencer Backman and Francisco Santos. In the second half, we give a brief survey on the theory of *cellular trees*, and try to provide partial results and evidence that suggest how our theory might shed some light on the subject.

CHAPTER 2

PRELIMINARIES

2.1 Graphs

We will assume basic definitions from standard references on graph theory such as [44]. All graphs we consider in this work are finite. Unless otherwise specified, we allow multi-edges but not loops. We will use n and m to denote the number of vertices and edges of a graph (whom will be clear from the context), respectively.

A *cycle* means simple cycle, i.e., a closed walk $v_0 - e_1 - v_1 - e_2 - \dots - v_k - e_k - v_0$ with all v_i 's pairwise distinct. Given a subset of vertices $U \subset V(G)$ and a vertex $v \in U$, $G[U]$ denotes the induced subgraph of G on U , and $\text{outdeg}_U(v)$ is the number of edges incident to v in which the other endpoint is outside U . Given two (disjoint) subsets U_1, U_2 of vertices, $G[U_1, U_2]$ is the subset of edges with one end in U_1 and the other end in U_2 . For a connected graph G , a subset of edges of the form $G[X, V(G) \setminus X]$ for some $\emptyset \neq X \subsetneq V(G)$ is a *cut*; an *isthmus* is an edge that forms a cut by itself. For a vertex $v \in V$, we often write $G[\{v\}, V \setminus \{v\}]$ as $\delta(v)$.

Let G be a connected graph. A *spanning tree* T is a connected spanning subgraph with no cycles; the set of spanning trees of G is denoted by $\mathcal{ST}(G)$. For every edge $e \notin T$, $T \cup \{e\}$ contains a unique cycle $C(T, e)$, it is the *fundamental cycle* of e with respect to T ; dually, for every edge $e \in T$, there is a unique cut $C^*(T, e)$ in G that does not contain any edges from $E(T) \setminus \{e\}$, it is the *fundamental cut* of e with respect to T . The *genus* g of G is the number of edges outside any spanning tree, which is $m - n + 1$.

An (*vertex-edge*) *incidence matrix* M_G is a matrix whose rows and columns are indexed by V and E , respectively, such that for each (non-loop) edge $e = uv$, the e -th column has a 1 in the u -th row and a -1 in the v -th row. The *Laplacian matrix* L_G is equal to $M_G M_G^T$, it is

independent of the choice of signs in M_G as it is easy to see that $L_G = D_G - A_G$, where D_G is a diagonal matrix whose (v, v) -entry is the degree of v , and A_G is the adjacency matrix of G . A *reduced Laplacian* $\widetilde{L}_G^{(i)}$ is L_G with the i -th row and the i -th column removed for some i ; $\widetilde{L}_G^{(i)}$ is equal to $\widetilde{M}_G^{(i)} \widetilde{M}_G^{(i)T}$, where $\widetilde{M}_G^{(i)}$ is M_G with the i -th row removed.

2.2 Divisors and Chip-firing

Definition 2.2.1 Let $G = (V, E)$ be a graph. A divisor on G is a function $D : V \rightarrow \mathbb{Z}$, which can be written as a formal sum $\sum_{v \in V} D(v) \cdot (v)$. D is an effective divisor if $D \geq 0$.

The set of all divisors on G is denoted by $\text{Div}(G)$; it has a natural group structure as the free abelian group generated by V . The degree $\deg(D)$ of a divisor D is $\sum_{v \in V} D(v)$, and the set of all degree k divisors is denoted by $\text{Div}^k(G)$.

Given a divisor D , we often say there are $D(v)$ chips at vertex v , and we can construct a divisor on G by adding or removing chips on vertices in various ways.

Definition 2.2.2 Let D be a divisor on G . A chip-firing move at a vertex v sends one chip away from v along each edge incident to v , hence v loses $\deg(v)$ chips and each neighbor w gains the number of edges between v and w many chips.

Two divisors D, D' are (linearly) equivalent, written as $D \sim D'$, if they differ by a sequence of chip-firing moves. The set of equivalence classes $\text{Div}^k(G) / \sim$ is denoted by $\text{Pic}^k(G)$. The group structure of $\text{Div}(G)$ restricts and descends to $\text{Pic}^0(G)$, which is called the (degree 0) Picard group of G .

It is not hard to see that if we identify divisors with vectors in \mathbb{Z}^n , then $D \sim D'$ if and only if $D - D' = L_G \mathbf{u}$ for some integral vector \mathbf{u} . In particular, $\text{Pic}^0(G) \cong \text{coker}_{\mathbb{Z}} \widetilde{L}_G$.

We introduce a few important classes of special divisors.

Definition 2.2.3 Let \mathcal{O} be an orientation of a graph G . The divisor $D_{\mathcal{O}}$ is $\sum_{v \in V(G)} (\text{indeg}_{\mathcal{O}}(v) - 1)(v)$. A divisor is orientable if it can be obtained in this way.

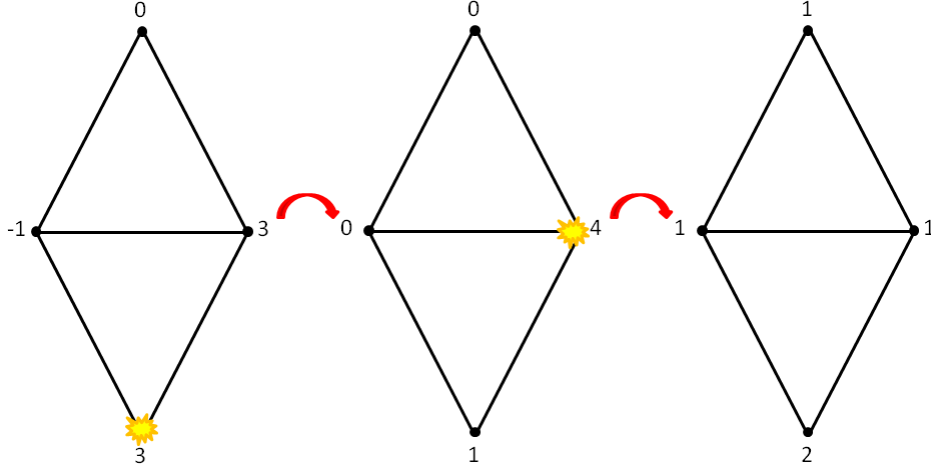


Figure 2.1: Example of chip-firing moves.

Fix $q \in V(G)$. An orientation \mathcal{O} is q -connected if any vertex of G can be reached from q via a directed path. A divisor is q -connected (known as q -orientable in [1]) if it is of the form $D_{\mathcal{O}}$ for some q -connected orientation \mathcal{O} .

Proposition 2.2.4 [1, Theorem 4.13] For any vertex q , the set of q -connected divisors form a set of representatives for the equivalence classes in $\text{Pic}^{g-1}(G)$.

Given a divisor D , the Euler characteristic $\chi_{D,G}$ sends each subset of vertices $S \subset V(G)$ to $\deg(D|_S) - g(G[S]) + 1$. We have the following criteria of orientable and q -connected divisors.

Proposition 2.2.5 [1, Theorem 4.8 and Theorem 4.11] A degree $g - 1$ divisor D is orientable if and only if $\chi_{D,G}(S) \geq 0$ for every non-empty $S \subset V(G)$. An orientable divisor D is q -connected if and only if $\chi_{D,G}(S) > 0$ for every non-empty $S \subset V(G) \setminus \{q\}$.

Definition 2.2.6 Fix a vertex $q \in V$, a divisor D is q -reduced (also known as superstable configurations [65] or G -parking functions [95, 113]) if

1. $D(v) \geq 0$ for each $v \neq q$, and
2. for any $U \subset V \setminus \{q\}$, there exists some $u \in U$ such that $D(u) < \text{outdeg}_U(u)$.

In chip-firing terms, (1) is to say that every vertex except possibly q is out of debt, while (2) is to say that firing any subset of V avoiding q will send some vertex into debt.

Proposition 2.2.7 [11, Proposition 3.1] For any integer k and vertex q , the set of degree k q -reduced divisors form a set of representatives for the equivalence classes in $\text{Pic}^k(G)$.

Dhar gave an algorithmic criterion for q -reduced divisors. Let D be a divisor that is effective outside q . Consider the following process: put $D(v)$ “firefighters” at each vertex $v \neq q$, where each firefighter can control fire from a single direction (i.e., an edge) leading into v . Start a fire at q and let the fire spread along the edges. As soon as the number of burnt edges incident to a vertex exceeds the number of firefighters there, the firefighters leave and the vertex will be burnt. We are interested in whether the process will burn the whole graph. A formal pseudocode is given as Algorithm 1.

```

Input: A divisor  $D \in \text{Div}(G)$ , and a vertex  $q \in V$ .
Output: A boolean value TRUE or FALSE.

if  $D(v) < 0$  for some  $v \neq q$  then
  | Output FALSE and Stop.
end
Set  $A_0 := V, v_0 := q$ . (Unburnt vertices, and the latest burnt vertex)
for  $1 \leq i \leq n - 1$  do
  |  $A_i := A_{i-1} \setminus v_{i-1}$ 
  | if  $\forall v \in A_i, D(v) \geq \text{outdeg}_{A_i}(v)$  then
  | | Output FALSE and Stop.
  | else
  | | Set  $v_i$  to be any vertex in  $A_i$  with  $D(v) < \text{outdeg}_{A_i}(v)$ 
  | end
end
Output TRUE.

```

Algorithm 1: Dhar’s Burning Algorithm

Proposition 2.2.8 [42] A divisor D is q -reduced if and only if Algorithm 1 returns TRUE.

2.3 Matroids and Oriented Matroids

2.3.1 Matroids

The concept of matroids was introduced by Whitney as an abstraction of independence in linear algebra and graph theory [110]. A comprehensive introduction on matroid theory is the book by Oxley [92].

Definition 2.3.1 A matroid $M = (E, \mathcal{I})$ consists of a non-empty finite set E (the ground set) and a collection \mathcal{I} of subsets of E that satisfies the following axioms:

1. $\emptyset \in \mathcal{I}$.
2. If $X \in \mathcal{I}$ and $Y \subset X$, then $Y \in \mathcal{I}$.
3. If $X, Y \in \mathcal{I}$ and $|X| < |Y|$, then there exists $y \in Y \setminus X$ such that $X \cup \{y\} \in \mathcal{I}$.

A subset of E is independent if it is in \mathcal{I} , it is dependent otherwise.

A *basis* is a maximal independent set with respect to inclusion; the set of bases of a matroid M is denoted by $\mathcal{B}(M)$. By the third axiom above, any two bases have the same size, which is denoted as the *rank* of the matroid; more generally, given a subset X of E , the rank $r(X)$ of X is the size of a maximal independent set contained in X . A *circuit* is a minimal dependent set; the set of circuits of a matroid M is denoted by $\mathcal{C}(M)$. The *closure* \overline{X} of a subset X is the maximal subset containing X such that $r(\overline{X}) = r(X)$, the closure is well-defined [92, Section 1.4].

Unless otherwise specified, we use m and r to denote the number of elements and the rank of a matroid (whom will be clear from the context), respectively.

Remark. It is possible to define a matroid using axiom systems based on each of the above notions [92, Chapter 1]. The phenomenon of having many equivalent axiom systems for matroids is known as *cryptomorphism* in the literature.

Two important operations in matroid theory are taking duals and taking minors, respectively.

Definition 2.3.2 Let $M = (E, \mathcal{I})$ be a matroid. The dual matroid M^* of M is a matroid on the same ground set E , where a subset Y is independent if and only if Y is disjoint from some basis of M . In particular, B is a basis of M if and only if its complement $B^* := E \setminus B$ is a basis of M^* . A circuit of M^* is said to be a cocircuit of M ; the set of cocircuits of a matroid M is denoted by $\mathcal{C}^*(M)$.

Let $X \subset E$ be a subset. The restriction $M|_X$ of M to X is the matroid $(X, \{I \in \mathcal{I} : I \subset X\})$. The deletion $M \setminus X$ of X from M is the matroid $M|_{E \setminus X}$. The contraction M/X of X from M is the matroid $(M^* \setminus X)^*$. A minor of M is a matroid obtained by a sequence of deletions and contractions from M . When $X = \{e\}$ is a singleton, we often write $X \setminus e$ and X/e instead of $X \setminus \{e\}$ and $X/\{e\}$.

Similar to the case of graphs, we have the following terminology in matroid theory.

Definition 2.3.3 An element of E is a loop if it is contained in no independent sets, or equivalently if it is a circuit by itself; dually, an element is an isthmus (coloop) if it is contained in every basis, or equivalently if it is a cocircuit by itself.

Given a basis B of M , the fundamental circuit $C(B, e)$ of an element $e \notin B$ with respect to B is the unique circuit contained in $B \cup \{e\}$; dually, the fundamental cocircuit $C^*(B, f)$ of an element $f \in B$ with respect to B is the unique cocircuit contained in $(E \setminus B) \cup \{e\}$.

Finally, we give some important examples of matroids.

Example 2.3.4 The uniform matroid $U_{r,m}$ is the matroid on m elements in which a subset X is independent if and only if $|X| \leq r$.

Example 2.3.5 The graphic matroid $M(G)$ of a finite graph $G = (V, E)$ is the matroid on E in which a subset $F \subset E$ is independent if and only if the induced subgraph (V, F) is a forest, i.e., if it contains no cycles.

If G is connected, then the circuits, cocircuits, and bases of $M(G)$ are (the edge sets of) the cycles, minimal cuts, and spanning trees of G , respectively. If G is furthermore a plane graph, then $M(G)^*$ is the graphic matroid of the dual graph G^* of G .

Example 2.3.6 The Fano matroid F_7 is the matroid on 7 elements illustrated by Figure 2.2, here a subset X is independent if and only if $|X| < 3$ or $|X| = 3$ and the three elements of X are not collinear.

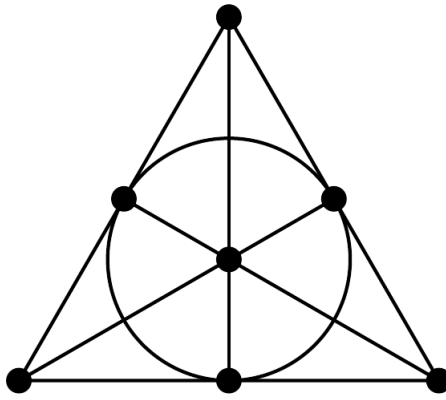


Figure 2.2: Fano matroid.

Example 2.3.7 Let A be a matrix over a field. $M(A)$ is the matroid whose ground set is the set of columns of A (as abstract elements), where a subset is independent if and only if the corresponding columns are linearly independent; in such case, we say $M(A)$ is represented by A . A matroid is representable if it can be represented by some matrix over some field. Every graphic matroid is representable as it is represented by the vertex-edge incidence matrix of the graph.

Essentially by definition, a subset is a circuit of $M(A)$ if and only if it is the support of some non-zero vector in the kernel of A that is minimal with respect to inclusion, and a subset is a cocircuit if and only if it is the support of some non-zero vector in the row space of A that is minimal with respect to inclusion.

2.3.2 Oriented Matroids

An *oriented matroid*, roughly speaking, is a matroid together with some “sign” data. The notion of oriented matroids was introduced independently by several groups of people, notably Robert Bland, Jon Folkman, Michel Las Vergnas, and Jim Lawrence [25, 51]. A comprehensive introduction on oriented matroid theory is the book by Björner, Las Vergnas, Sturmfels, White, and Ziegler [23].

We make precise the definition of “sign data”. Given a non-empty, finite ground set E , a *signed subset* of E is a map $X : E \rightarrow \{+, 0, -\}$ that assigns signs to its elements; we define a total ordering of $\{+, 0, -\}$ as $+ > 0 > -$. The *support* \underline{X} of a signed subset X is $\{e \in E : X(e) \neq 0\}$. Similarly, we define X^+ and X^- to be $\{e \in E : X(e) = +\}$ and $\{e \in E : X(e) = -\}$, respectively. Now we have the following definition.

Definition 2.3.8 *An oriented matroid $M = (E, \mathcal{C})$ consists of a non-empty, finite set E and a collection \mathcal{C} of signed subsets that satisfies the following axioms:*

1. $\emptyset \notin \mathcal{C}$.
2. If $C \in \mathcal{C}$, then $-C \in \mathcal{C}$.
3. For all $C_1, C_2 \in \mathcal{C}$, if $\underline{C_1} \subset \underline{C_2}$, then either $C_1 = C_2$ or $C_1 = -C_2$.
4. For all $C_1, C_2 \in \mathcal{C}$, $e \in C_1^+ \cap C_2^-$ and $f \in (C_1^+ \setminus C_2^-) \cup (C_1^- \setminus C_2^+)$, there is a $C_3 \in \mathcal{C}$ such that $C_3^+ \subset (C_1^+ \cup C_2^+) \setminus \{e\}$, $C_3^- \subset (C_1^- \cup C_2^-) \setminus \{e\}$, and $f \in C_3$.

A signed subset in \mathcal{C} is said to be a signed circuit.

By [92, Theorem 1.1.4], $\{\underline{C} : C \in \mathcal{C}\}$ is the collection of circuits of a matroid \underline{M} with ground set E , called the *underlying matroid* of M . Two signed subsets X, Y are *orthogonal* if either $\underline{X} \cap \underline{Y} = \emptyset$, or there exist $e, f \in E$ such that $X(e) = Y(e) \neq 0$ and $X(f) = -Y(f) \neq 0$. There is a canonical way to assign signings to the cocircuits of \underline{M} so that they are orthogonal to all signed circuits of M [23, Theorem 3.4.1], such signings

of cocircuits are the *signed cocircuits* of M . The *dual matroid* M^* of M is the oriented matroid on the same ground set whose signed circuits are the signed cocircuits of M .

Let X, Y be two signed subsets. Their *composition* $X \circ Y$ is the signed subset given by $X \circ Y(e) = X(e)$ if $X(e) \neq 0$, and $X \circ Y(e) = Y(e)$ otherwise. A composition of signed circuits is a (*signed*) *vector*, and a composition of signed cocircuits is a (*signed*) *covector*. Moreover, we say X and Y are *conformal* if there are no elements e such that $X(e) = -Y(e) \neq 0$.

We have other oriented matroidal operations as in ordinary matroids: given an oriented matroid $M = (E, \mathcal{C})$ and a subset $X \subset E$, the *restriction* $M|_X$ is the oriented matroid $(X, \{C \in \mathcal{C} : \underline{C} \subset X\})$. The *deletion* $M \setminus X$ is the restriction of M to $E \setminus X$. The *contraction* M/X is the oriented matroid on ground set $E \setminus X$ whose signed circuits are the support minimal signed subsets of $\{C|_{E \setminus X} : C \in \mathcal{C}\}$, here $C|_X$ is $C : E \rightarrow \{+, -, 0\}$ restricted to X . A *minor* of M is an oriented matroid obtained by a sequence of deletions and contractions from M .

Given an oriented matroid $M = (E, \mathcal{C})$ and a subset $X \subset E$, the *reorientation* of M with respect to X negates the signs of the coordinates over X for each signed circuit. An *orientation* of M is a map from E to $\{+, -\}$. In some literature, an orientation is identified with a reorientation of the oriented matroid, but since orientations have intuitive meaning in the case of graphs as directed graphs, we will stick with our terminology.

Every matrix A over an ordered field (which can be assumed to be the field of real numbers [23, Proposition 8.4.1]) gives an oriented matroid $M(A)$, where the ground set is the columns of A , and the signed circuits (resp. signed cocircuits) are the sign patterns of the support minimal non-zero vectors in the kernel (resp. row space) of A . Such oriented matroids are said to be *realizable*, and we say A is a *realization* of M if $M = M(A)$.

Remark. On one hand, not every matroid is orientable (cf. [23, Example 6.6.2]). On the other hand, a matroid can have multiple oriented structures that are not equivalent

even up to reorientation (cf. [23, Section 1.5]). Furthermore, the same oriented structure can be realized by very different matrices, as an example, any matrix of the form $\begin{pmatrix} t_1 & t_2 & t_3 & t_4 \\ at_1 & bt_2 & ct_3 & dt_4 \end{pmatrix}$, $t_1, t_2, t_3, t_4 > 0, a > b > c > d > 0$ realizes the same oriented structure of $U_{2,4}$ (the *four-point line*).

We state a technical but elementary lemma that is useful when studying realizable oriented matroids. In abstract oriented matroid terms, it is to say that every signed vector of an oriented matroid is a conformational composition of signed circuits.

Lemma 2.3.9 [115, Lemma 6.7] *Let $\mathbf{u} \in \mathbb{R}^E$ be a vector in $\ker(A)$. Then \mathbf{u} can be written as $\sum \mathbf{v}_C$ such that:*

1. *Each \mathbf{v}_C is from $\ker(A)$.*
2. *The sign pattern C of each \mathbf{v}_C is a signed circuit of $M(A)$.*
3. *The support of each C is contained in the support of \mathbf{u} .*
4. *For each $e \in \underline{C}$, the sign of e in C and the sign of the e -th coordinate of \mathbf{u} agree.*

2.3.3 Regular Matroids

A matrix A over \mathbb{R} is *totally unimodular* if the determinant of every square submatrix is either 0, 1 or -1 . In particular, every totally unimodular matrix is an integer matrix. We will consider a special class of matroids represented by such matrices.

Definition 2.3.10 *A matroid is regular if it can be represented by a totally unimodular matrix.*

Since the vertex-edge incidence matrix of a graph is totally unimodular [92, Lemma 5.1.4], every graphic matroid is regular.

From the definition, every regular matroid is orientable. The family of regular matroids is closed under taking duals and minors [92, Proposition 2.2.22, Proposition 3.2.5]. Regular

matroids have many equivalent characterizations of different favors. We list a few important ones in the following theorem.

Theorem 2.3.11 *The following statements are equivalent for a matroid M .*

1. M is regular.
2. M is representable over every field.
3. M has no minor isomorphic to any of $U_{2,4}$, F_7 or F_7^* .
4. M is binary and orientable.

PROOF: The equivalence of (1), (2), and (3) is part of [109, Theorem 3.1.1]. (2) implying (4) is trivial, while the converse is [23, Theorem 7.9.3]. \square

A more complete list can be found in [109, Theorem 3.1.1], see also Seymour's decomposition theorem [92, Theorem 13.1.1].

Using Theorem 2.3.11, we have the following corollary.

Proposition 2.3.12 *An oriented matroid is regular if and only if the underlying matroid has no $U_{2,4}$ -minor.*

PROOF: By Theorem 2.3.11, an oriented matroid is regular if and only if it is binary. By [92, Theorem 6.5.4], a matroid is binary if and only if it has no $U_{2,4}$ -minor. \square

Regular matroids are rigid in the sense that the underlying matroid completely determines its oriented structures and realizations up to trivial differences.

Proposition 2.3.13 *The oriented structures on a regular matroid differ only by reorientations. Furthermore, if A, A' are two $r \times m$ totally unimodular matrices of full row rank that realize the same regular matroid, then there exists a unimodular integer matrix P such that A' equals PA with some (possibly none) columns negated. In particular, P preserves volume in \mathbb{R}^r and induces a bijection of lattice points in \mathbb{Z}^r .*

PROOF: The first statement is [23, Corollary 7.9.4]. The second statement can be found in [105]. \square

The significance of Proposition 2.3.13 is that, while the definitions of several objects associated to a regular matroid involve choosing an oriented structure or a totally unimodular realization, these objects (or at least most of their properties) are intrinsic and independent of such choices.

We state a useful lemma for regular matroids. Note that total unimodularity is crucial in the proof, and the lemma is not true for general oriented matroids. The failure of the lemma is one of the main reasons that some of our definitions and results cannot be generalized beyond regular matroids.

Lemma 2.3.14 *Let M be a regular matroid and let A be an arbitrary totally unimodular matrix realizing M . Then any support minimal non-zero vector \mathbf{u} in the kernel of A is a positive multiple of some $\{0, 1, -1\}$ -vector $\hat{\mathbf{u}}$, which can be interpreted as a signed circuit directly by identifying $1, 0, -1$ with $+, 0, -$, respectively. Conversely, under such identification, every signed circuit C can be thought as an element in $\ker(A) \subset \mathbb{R}^E$.*

Furthermore, if \mathbf{v} is a vector in $\ker(A)$ whose coordinates are $0, 1, -1$, then the \mathbf{v} can be written as a sum of disjoint signed circuits.

PROOF: By restricting and negating the columns of A if necessary, we may assume that all coordinates of \mathbf{u} are positive. Let \mathbf{a} be the first column of A (which is a $\{0, 1, -1\}$ -vector) and let A' be the rest of the matrix. The columns of A' are linearly independent as the support of \mathbf{u} is a circuit of M , so we can pick a non-singular maximal square submatrix \tilde{A} of A' . Denote by $\tilde{\mathbf{a}}$ the corresponding restriction of \mathbf{a} , the unique solution of $\tilde{A}\mathbf{x} = -\tilde{\mathbf{a}}$ is then \mathbf{u}/u_1 excluding the first coordinate. By Cramer's rule and the total unimodularity of \tilde{A} , all coordinates of \mathbf{u}/u_1 are $0, 1$, or -1 . The converse is trivial as every signed circuit is by definition the sign pattern of some support minimal non-zero vector of $\ker(A)$.

For the second part of the lemma, we again assume that \mathbf{v} is a $\{0, 1\}$ -vector by negating columns of A . By Lemma 2.3.9, there exists a non-negative signed circuit C whose support \underline{C} is contained in the support X of \mathbf{v} . By the first half of this lemma, $\mathbf{v} - C$ is a $\{0, 1\}$ -vector in $\ker(A)$ whose support is $X \setminus \underline{C}$. By induction on the size of support, $\mathbf{v} - C$ can be written as the sum of disjoint signed circuits, and these circuits are disjoint from C . \square

2.3.4 Circuit and Cocircuit Lattices

From Section 2.3.3, signed circuits and signed cocircuits of a regular matroid M can be identified as $\{0, 1, -1\}$ -vectors. This leads to the following definitions.

Definition 2.3.15 *Let M be a regular oriented matroid on ground set E . The lattice of 1-chains $C_1(M) \cong \mathbb{Z}^E$ is the free \mathbb{Z} -module generated by the elements of E . The lattice of integer flows $Z_1(M)$ is the sublattice of $C_1(M)$ generated by all signed circuits of M ; the lattice of integer cuts $B_1(M)$ is the sublattice of $C_1(M)$ generated by all signed cocircuits of M .*

Embed $C_1(M)$ into $C_1(M) \otimes \mathbb{R} \cong \mathbb{R}^E$. The affine spans of $Z_1(M)$ and $B_1(M)$ are the circuit space $V(M)$ and the cocircuit space $V^(M)$ of M , respectively.*

We collect a few properties of these objects.

Proposition 2.3.16 *Let M be a regular oriented matroid. Then*

1. *For any totally unimodular matrix A realizing M , the circuit space and cocircuit space are the kernel and the row space of A , respectively.*
2. *$V(M)$ is the orthogonal complement of $V^*(M)$; $V(M)$ and $V^*(M)$ are of dimension $m - r$ and r , respectively.*
3. *$C_1(M) \cap V(M) = Z_1(M)$, $C_1(M) \cap V^*(M) = B_1(M)$.*

4. For any basis B of M , the set of fundamental circuits $\{C(B, e) : e \notin B\}$ with respect to B forms a \mathbb{Z} -basis (resp. \mathbb{R} -basis) of $Z_1(M)$ (resp. $V(M)$); dually, the set of fundamental cocircuits $\{C^*(B, e) : e \in B\}$ forms a \mathbb{Z} -basis (resp. \mathbb{R} -basis) of $B_1(M)$ (resp. $V^*(M)$).

PROOF: (1), (2) follow from Lemma 2.3.14; (4) is from [105, Section 2.3]. For (3), let $\mathbf{z} \in C_1(M) \cap V(M)$ be an element. By (1), \mathbf{z} is in the kernel of A , so by Lemma 2.3.9 and 2.3.14, there exists some signed circuit C of M such that the support of C is contained in the support of \mathbf{z} . Pick any $e \in \underline{C}$, then $\mathbf{z}' := \mathbf{z} - (z_e/C(e))C$ is an element in $C_1(M) \cap V(M)$ whose support is strictly contained in the support of \mathbf{z} . By induction on the size of support, we can write \mathbf{z} as a linear combination of signed circuits over \mathbb{Z} , that is, $\mathbf{z} \in Z_1(M)$. The converse inclusion is obvious, and the dual statement for $B_1(M)$ is similar. \square

Remark. $Z_1(M)$ and $B_1(M)$ can be defined for any oriented matroid, but they will be less well-behaved because the lack of property (1). In fact, property (2) characterizes regular matroids among oriented matroids, cf. [89], [91, Section 4.1].

2.4 Jacobian Groups

The Jacobian is a finite abelian group canonically associated to a regular matroid.

Definition 2.4.1 *Let M be a regular matroid. The Jacobian $\text{Jac}(M)$ of M is the group*

$$\frac{C_1(M)}{Z_1(M) \oplus B_1(M)}.$$

By [87] (or more directly, Proposition 2.3.13), $\text{Jac}(M)$ does not depend on the choice of the oriented structure of M .

Proposition 2.4.2 *Let M be a regular matroid and fix a full row rank, totally unimodular representation A of M . Then $[\gamma] \mapsto [A\gamma]$ is a group isomorphism between $\frac{C_1(M)}{Z_1(M) \oplus B_1(M)}$ and $\text{coker}_{\mathbb{Z}}(AA^T)$, here $Z_1(M)$, $B_1(M)$ are defined using the oriented structure of M realized by A .*

PROOF: The map is well-defined because $A(Z_1(M) \oplus B_1(M)) = A(B_1(M)) = A(\text{Col}_{\mathbb{Z}} A^T) = \text{Col}_{\mathbb{Z}} AA^T$, the equality also shows the map is injective. It is surjective because $A\gamma = \mathbf{b}$ has a solution in $C_1(M)$ for every $b \in \mathbb{Z}^r$, using the total unimodularity of A . \square

Corollary 2.4.3 *For a graph G , $\text{Jac}(M(G)) \cong \text{Pic}^0(G)$.*

We have following generalization of Kirchhoff's Matrix–Tree theorem, which is usually stated as the determinant of any reduced Laplacian (or product of the non-zero eigenvalues of the Laplacian) equals the number of spanning trees.

Theorem 2.4.4 *The cardinality of $\text{Jac}(M)$ equals the number of bases of M . In particular, $|\text{Pic}^k(G)|$ equals the number of spanning trees of G for any k .*

PROOF: By Proposition 2.4.2, it suffices to compute $|\text{coker}_{\mathbb{Z}}(AA^T)|$. By Cauchy–Binet formula, $|\text{coker}_{\mathbb{Z}}(AA^T)| = \det(AA^T) = \sum_{I \subset E, |I|=r} \det(A|_I)^2$. But since A is totally unimodular, $\det(A|_I)^2$ equals 1 if I is a basis, and 0 otherwise, so the sum counts the number of bases of M . \square

2.5 Tutte Polynomials

The Tutte polynomial is a bivariate polynomial invariant associated to every matroid, which satisfies a deletion-contraction recurrence relation [106]. For the rest of this section, M will be a matroid on ground set E .

Definition 2.5.1 *The Tutte polynomial $T_M(x, y)$ of M is defined recursively as follows:*

- *If $e \in E$ is neither a loop nor an isthmus, then $T_M(x, y) = T_{M \setminus e}(x, y) + T_{M/e}(x, y)$, otherwise*
- *$T_M(x, y) = x^i y^j$ if M consists only of i isthmus and j loops.*

Tutte polynomials have many equivalent definitions and descriptions, including a definition using the rank function and a definition using *basis activities*.

Proposition 2.5.2 [28, Lemma 6.2.1] Denote by $r : 2^E \rightarrow \mathbb{Z}$ the rank function of M . Then

$$T_M(x, y) = \sum_{A \subseteq E} (x - 1)^{r(M) - r(A)} (y - 1)^{|A| - r(A)}.$$

Definition 2.5.3 Let $<$ be a total ordering of E . Given a basis B of M . An element $e \notin B$ (resp. $e \in B$) is externally active (resp. internally active) if e is the smallest element in its fundamental circuit (resp. fundamental cocircuit) with respect to B . The number of externally active (resp. internally active) elements with respect to B is the external activity $\epsilon(T)$ (resp. internal activity $\iota(T)$) of T .

Proposition 2.5.4 [28, Chapter 6.6A] For any total ordering of E , we have

$$T_M(x, y) = \sum_{B \in \mathcal{B}(M)} x^{\iota(B)} y^{\epsilon(B)}.$$

The Tutte polynomial is universal for deletion-contraction invariants [28, Theorem 6.2.2]. In particular, many important enumerative quantities of a matroid can be expressed as special evaluations of the Tutte polynomial. Here are a few basic ones.

Proposition 2.5.5 [28, Proposition 6.2.11] Let M be a matroid. Then

1. $T_M(1, 1)$ equals the number of bases of M .
2. $T_M(2, 1)$ equals the number of independent sets of M .
3. $T_M(1, 2)$ equals the number of spanning sets of M , that is, subsets whose closure is the all of E .

Other prominent specializations of the Tutte polynomial include the *chromatic polynomial* and *flow polynomial* of a matroid [109, Chapter 7], see also [27, Chapter X] for applications in statistical physics and knot theory; some connections with oriented matroids will be discussed in Section 4.3 and 5.5.

CHAPTER 3

COMBINATORICS OF BREAK DIVISORS AND THE ABKS DECOMPOSITION

3.1 Definitions

3.1.1 Metric Graphs

A *metric graph* Γ is a compact connected metric space such that every point has a neighborhood isometric to a star-shaped set. Every metric graph can be constructed in the following way: start with a weighted graph $(G = (V, E), w : E \rightarrow \mathbb{R}_{>0})$, associate with each edge e a closed line segment L_e of length $w(e)$, and identify the endpoints of the L_e 's according to the graph structure in the obvious way. A (weighted) graph G that yields a metric graph Γ is said to be a *model* of Γ . The *genus* of Γ is the genus of any model of Γ .

A *divisor* on Γ is a function $D : \Gamma \rightarrow \mathbb{Z}$ of finite support, or equivalently a finite formal sum of points of Γ over \mathbb{Z} . A *cut* Z of Γ is a disjoint union of equal length open line segments such that $\Gamma \setminus Z$ is a disjoint union of two connected closed sets X, Y . A *chip-firing* move chooses a cut $Z = Z_1 \sqcup \dots \sqcup Z_k$ of the metric graph and for every line segment Z_i , moves a chip from the endpoint of $\overline{Z_i}$ in X to the endpoint of $\overline{Z_i}$ in Y . We can define notions such as degree, effective divisors, (linearly) equivalence, $\text{Div}(\Gamma)$, $\text{Div}^k(\Gamma)$, and $\text{Pic}^k(\Gamma)$ analogous to their discrete counterparts.

A divisor on a metric graph Γ is *integral* with respect to a model G if the support of D consists of vertices of G , and in such case the divisor can be considered as a divisor on G . Given a divisor $D \in \text{Div}(\Gamma)$ (and a point $q \in \Gamma$), by choosing a model G of Γ such that $V(G)$ includes the support of D (and q), we say D is *orientable* (resp. *q -connected*, *q -reduced*) if D is so when considered as a divisor on G . The criteria in Proposition 2.2.5 and 2.2.8 can be applied after such reduction.

Finally, we define the *rank* of a divisor in tropical Brill-Noether theory [9, 33]. A

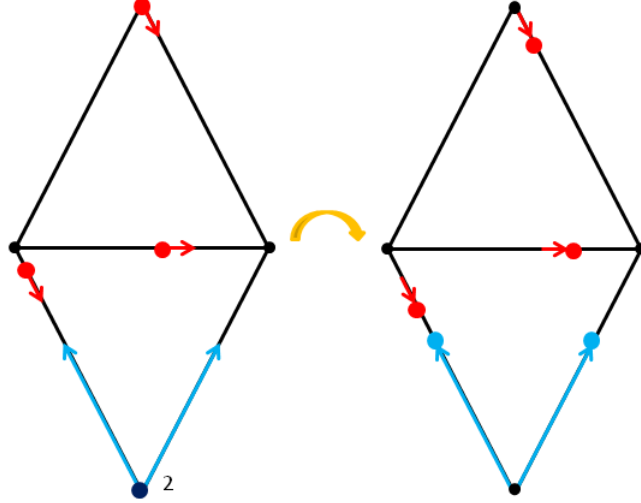


Figure 3.1: Examples of chip-firing on a metric graph.

divisor $D \in \text{Div}(\Gamma)$ is of rank -1 if it is not linearly equivalent to any effective divisors. Otherwise, the rank of D is the largest number r such that no matter how one takes away r chips from D , the new divisor is linearly equivalent to some effective divisor. The *Brill-Noether locus* $W_d^r(\Gamma)$ is the set of divisor classes in $\text{Pic}^d(\Gamma)$ whose rank is at least r ; a geometric structure will be given to $W_d^r(\Gamma)$, as a subset of $\text{Pic}^d(\Gamma)$, in Section 3.1.3. The *Brill-Noether number* $\rho(g, r, d) := g - (r + 1)(g - d + r)$ is the expected dimension of $W_d^r(X)$ for a genus g algebraic curve X in the classical algebro-geometric setting [60].

3.1.2 Break Divisors

Definition 3.1.1 A divisor D of a connected graph G is a break divisor if it can be obtained from the following procedure: choose a spanning tree T of G , and for every edge $f \notin T$, pick an orientation of f and add a chip at the head of f . The set of break divisors of G is denoted by $\mathcal{BD}(G)$.

We collect a few important properties of break divisors.

Proposition 3.1.2 Let G be a connected graph. Then

1. Every break divisor is effective and of degree g .

2. A divisor D is break if and only if $D - (q)$ is a q -connected divisor for every vertex q of G .
3. $\mathcal{BD}(G)$ is a set of representatives for the divisor classes of $\text{Pic}^g(G)$.
4. The number of break divisors of G equals the number of spanning trees of G .

PROOF: (1) is trivial from definition. (2) and (3) are Lemma 3.3 and Theorem 1.3 of [1], respectively. (4) is a simply corollary of (3) and Theorem 2.4.4. \square

We mention another characterization of break divisors using Euler characteristic.

Lemma 3.1.3 *A degree g divisor D is break if and only if $\chi_{D,G}(S) > 0$ for every non-empty $S \subset V(G)$.*

PROOF: Fix an arbitrary vertex q . D is break if and only if $D - (q)$ is q -connected, if and only if, by Proposition 2.2.5, (1) for every non-empty $S \subset V(G) \setminus \{q\}$, $\chi_{D,G}(S) = \chi_{D-(q),G}(S) > 0$ and (2) for every $q \in S \subset V(G)$, $\chi_{D,G}(S) = \chi_{D-(q),G}(S) + 1 > 0$. \square

The following definition generalizes the notion of break divisors to metric graphs.

Definition 3.1.4 *Let Γ be a metric graph and fix an arbitrary model G of Γ . A divisor D of Γ is a break divisor if it can be obtained from the following procedure: choose a spanning tree T of G , and for every edge $f \notin T$, put a chip inside the interval of Γ corresponding to f (possibly at an endpoint). In such case we say the break divisor is associated to T (such T is in general not unique).*

Theorem 3.1.5 *[1, Theorem 1.1] (also [88]) Break divisors on Γ form a set of representatives for the equivalence classes in $\text{Pic}^g(\Gamma)$.*

3.1.3 Tropical Jacobians and the ABKS Decomposition

The Picard group $\text{Pic}^0(\Gamma)$, and hence any $\text{Pic}^k(\Gamma)$, has a natural structure as a g -dimensional real torus. We give an informal description here; a rigorous treatment is given in [10, 88]. The theory here is can be thought as a geometric version of Section 2.3.4 and 2.4.

Fix a model $G = (V, E, w)$ for Γ , as well as an arbitrary *reference orientation* for the edges of G . The *real edge space* $C_1(G; \mathbb{R})$ of G is the m -dimensional vector space with E as a basis over \mathbb{R} . This space is equipped with an *inner product* that extends $e_i \cdot e_j = \delta_{i,j} w(e_i)$ bilinearly.

The *real cycle space* $H_1(G; \mathbb{R})$ is the g -dimensional subspace of $C_1(G; \mathbb{R})$ spanned by all cycles of G ; denote by $\pi : C_1(G; \mathbb{R}) \rightarrow H_1(G; \mathbb{R})$ the orthogonal projection. By Proposition 2.3.16, the projections of k edges $f_1, \dots, f_k \in G$ onto $H_1(G; \mathbb{R})$ are linearly independent if and only if the set of edges does not contain a cut, if and only if $G \setminus \{f_1, \dots, f_k\}$ is connected; hence $\pi(e_1), \dots, \pi(e_g)$ form a basis of $H_1(G; \mathbb{R})$ exactly when e_1, \dots, e_g is the complement of some spanning tree.

The *integral cycle space* $H_1(G; \mathbb{Z})$ is the g -dimensional lattice in $H_1(G; \mathbb{R})$ consisting of integral combinations of cycles of G . For any spanning tree T of G , the set of fundamental cycles $\{C(T, e) : e \notin T\}$ forms a \mathbb{R} -basis (resp. \mathbb{Z} -basis) of $H_1(G; \mathbb{R})$ (resp. $H_1(G; \mathbb{Z})$), cf. Proposition 2.3.16 as well as [21]. The *tropical Jacobian* $\text{Jac}(\Gamma)$ is the g -dimensional real torus $H_1(G; \mathbb{R})/H_1(G; \mathbb{Z})$ with the induced inner product.

Now we define a bijection between $\text{Pic}^g(\Gamma)$ and $\text{Jac}(\Gamma)$. Fix a vertex q of G . From each divisor class in $\text{Pic}^g(\Gamma)$, pick any effective divisor $D = (p_1) + \dots + (p_g)$ (by Theorem 3.1.5, we can always choose a break divisor). For each p_i , choose a path γ_i from q to p_i and interpret γ_i as an element of $C_1(G; \mathbb{R})$. Then the image of $[D]$ in $\text{Jac}(\Gamma)$ is $[\pi(\gamma_1 + \dots + \gamma_g)] \in H_1(G; \mathbb{R})/H_1(G; \mathbb{Z}) = \text{Jac}(\Gamma)$. Different choices of q produce the same bijection up to translations in the universal cover of $\text{Jac}(\Gamma)$ only, so essentially the map is independent of q , and from now on we will often abuse notations and identify $\text{Pic}^g(\Gamma)$ and $\text{Jac}(\Gamma)$.

We also note that the map above can be used to define a map from $\text{Div}_+^g(\Gamma)$, the space of effective degree g divisors, to $\text{Jac}(\Gamma)$.

Example 3.1.6 In Figure 3.2, take $\{e_1, e_3, e_5\}$ as the spanning tree associated to the break

divisor. The cycle space is spanned by $C_1 := e_1 - e_2 + e_5$, $C_2 := e_3 - e_4 + e_5$. Take the paths $\gamma_1 = e_1, \gamma_2 = e_2 - \frac{1}{3}e_4$ from q to p_1, p_2 . The image of $(p_1) + (p_2)$ in $\text{Jac}(\Gamma)$ is $[\pi((e_1) + (e_2 - \frac{1}{3}e_4))] = [\frac{-1}{24}C_1 + \frac{1}{8}C_2] = [\frac{23}{24}C_1 + \frac{1}{8}C_2]$.

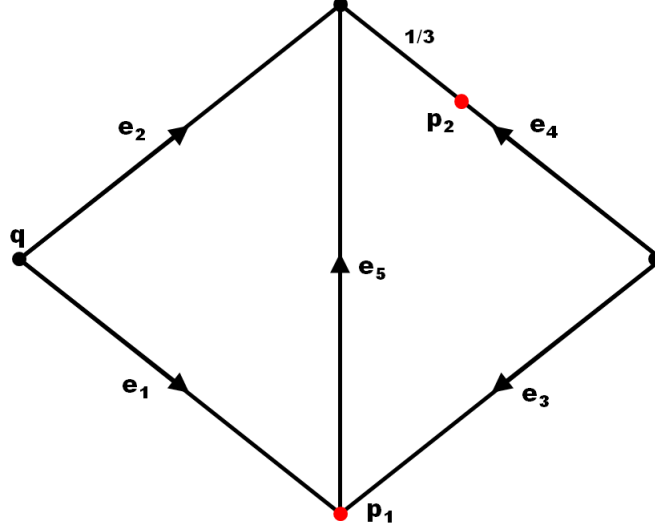


Figure 3.2: A fixed model for a metric graph Γ and a break divisor $(p_1) + (p_2)$ on it.

Next we study how the image of a break divisor changes when it is perturbed by a small amount; for simplicity, we assume $w \equiv 1$. Let T be a spanning tree and let the (unit length) edges not in T be $e_1 = \overline{u_1 v_1}, \dots, e_g = \overline{u_g v_g}$. Let $D = (p_1) + \dots + (p_g)$ be a break divisor where $p_i \in L_{e_i}$, say the segment $\overline{u_j p_j}$ is of length $0 \leq \theta_j < 1$. Suppose we perturb the chip at p_j and move it to $p'_j \in L_{e_i}$ where $\overline{u_j p'_j}$ is of length $\theta_j < \theta'_j \leq 1$. Then the image of $[(p_1) + \dots + (p'_j) + \dots + (p_g)]$ is $[(\theta'_j - \theta_j)\pi(e_j)]$ away from the image of $[D]$.

With T and $e_1 = \overline{u_1 v_1}, \dots, e_g = \overline{u_g v_g} \notin T$ as above, $D = (u_1) + \dots + (u_g)$ is a break divisor associated to T , and the image of every break divisor associated to T is equal to $[D] + [\sum_{i=1}^g \theta_i \pi(e_i)]$ for some $0 < \theta_i < 1$. Therefore the image of all such break divisors in $\text{Jac}(\Gamma)$, denoted by \mathcal{C}_T , is the image of (a translation of) the parallelotope $\{\sum_{i=1}^g \theta_i \pi(e_i) : 0 \leq \theta_i \leq 1\} \subset H_1(G, \mathbb{R})$ in $\text{Jac}(\Gamma)$. The tropical Jacobian is the union of \mathcal{C}_T 's as T runs through all spanning trees of G . As two distinct cells are disjoint except possibly at the boundary, they give a polyhedral decomposition of $\text{Jac}(\Gamma)$.

Definition 3.1.7 *The ABKS decomposition of $\text{Jac}(\Gamma)$, with respect to a model G of Γ , is the polyhedral decomposition described above.*

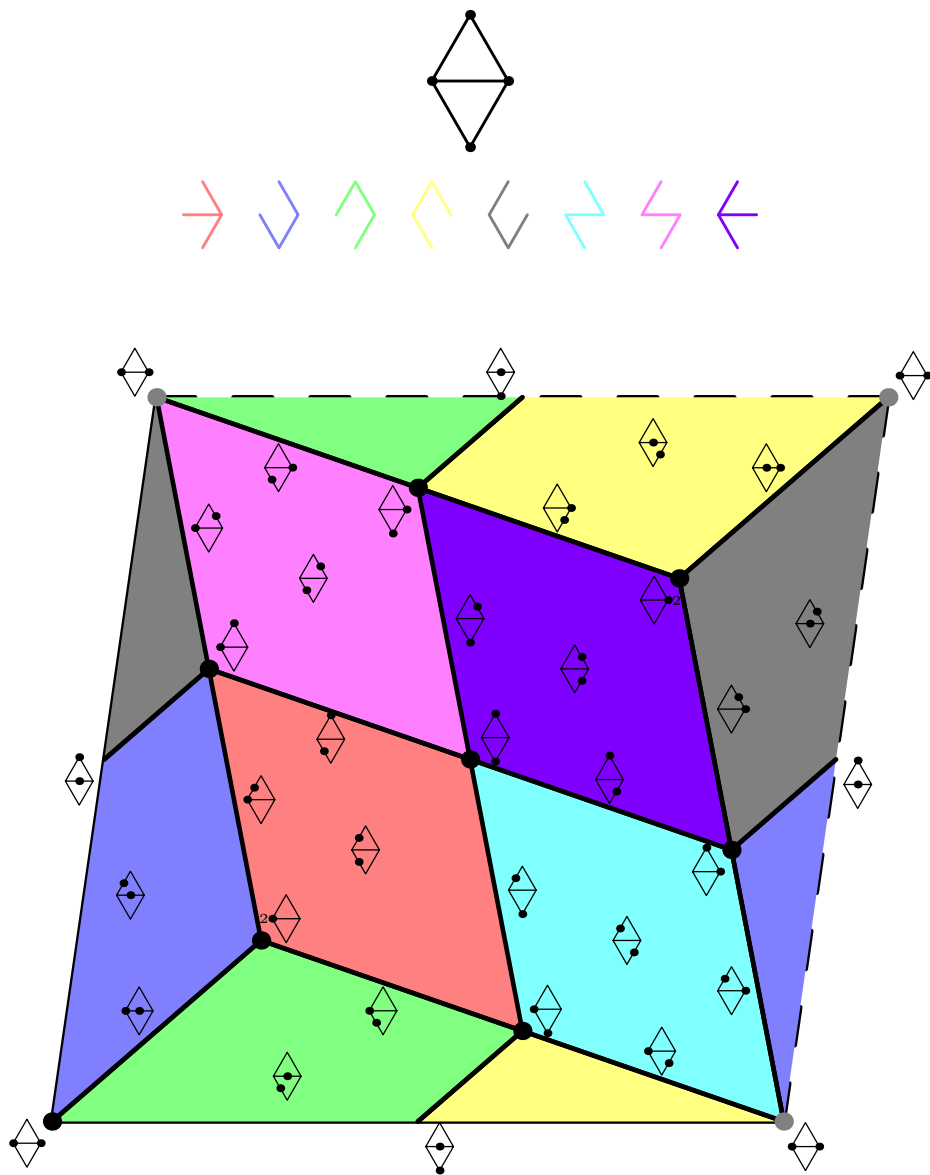


Figure 3.3: The ABKS decomposition of $\text{Pic}^2(\Gamma)$ ([1, Figure 1]), using the model chosen in Figure 3.2.

3.2 Checking Break Divisors by Matroid Intersection

The network flow based algorithm of Backman [4] gives an efficient method to check whether a given divisor is break. Here we give an alternative algorithmic approach to this problem using matroid intersection. Despite a slightly worse runtime, such point of view provides a polyhedral picture of the problem.

For an effective degree g divisor $D = (v_1) + \dots + (v_g)$, construct the bipartite graph G_D whose vertex set is $E \sqcup \{1, 2, \dots, g\}$ and $e \in E$ is adjacent to i precisely when e is incident to v_i in G . Such bipartite graph corresponds to a *transversal matroid*, in which $S \subset E$ is independent if and only if there exists a matching of G_D saturating S [92, Section 1.6].

Proposition 3.2.1 *An effective degree g divisor D is break if and only if some common independent set in the cographic matroid of G and the transversal matroid of G_D has cardinality g .*

PROOF: This follows almost directly from the definition of break divisors. If D is break, pick a spanning tree T and an orientation of edges outside T that induce D , then $S := E(G) \setminus E(T)$ is independent in the cographic matroid of G , and the orientation of edges in S specifies a matching of G_D by matching $e \in S$ to i if e is pointing to v_i (in general v_i 's are not distinct, but if e_1, \dots, e_k are both pointing to the same v , then we simply pick any matching between these edges and the k values in $\{1, 2, \dots, g\}$ that correspond to v), so S is independent in the transversal matroid as well. Conversely, if such common independent set S of cardinality g exists, then $E(G) \setminus S$ is a spanning tree of G and any matching in G_D saturating S specifies a way to orient edges in S that induces D . \square

Since efficient algorithms are known for solving the maximum common independent set problem for the intersection of two matroids [80], our approach yields a polynomial-time algorithm for checking whether a divisor is break. Geometrically speaking, Proposition 3.2.1 says a divisor is break if and only if the transversal matroid polytope corresponding to G_D has a non-empty intersection with the cographic matroid polytope of G .

3.3 Some Geometry of $\text{Div}_+^g(\Gamma) \rightarrow \text{Jac}(\Gamma)$

We show that the break divisor section is the unique section of the map $\text{Div}_+^g(\Gamma) \rightarrow \text{Jac}(\Gamma)$. The statement itself is not too hard to see from the polyhedral picture, as both $\text{Div}_+^g(\Gamma)$ and $\text{Jac}(\Gamma)$ are g -dimensional polyhedral complexes in which every fiber (as a tropical linear system [63]) is connected. However, our approaches provide connections with other topics.

We first give a proof that uses the combinatorics of tropical divisors. A more general version of the following lemma can be found in [88, Theorem 8.5], proven using sophisticated arguments, here we give a more intuitive proof to the special case.

Lemma 3.3.1 *Let D be a break divisor outside the $(g - 1)$ -skeleton of the ABKS decomposition (with respect to any model G). Then D is q -reduced for every $q \in \Gamma$.*

PROOF: We run Dhar’s burning algorithm and start a fire at q . Pick an arbitrary edge e that contains q . It has at most one interior chip (or “fireman”) and no chips at endpoints, so the fire can burn through at least one of the endpoints. From there the fire can burn through the remaining vertices of G using the (unique) spanning tree associated to D . Now for each edge of G , both of its endpoints are burnt while it has at most one interior chip, so the whole edge will be burnt. \square

Proposition 3.3.2 *$\text{Div}_+^g(\Gamma) \rightarrow \text{Jac}(\Gamma)$ has a unique continuous section.*

PROOF: In [1, 88], it was proven that the set of break divisors is a continuous section to $\text{Div}_+^g(\Gamma) \rightarrow \text{Jac}(\Gamma)$. For uniqueness, notice that if D is a generic break divisor in the sense of Lemma 3.3.1, then D is the unique effective divisor linearly equivalent to itself, so any continuous section of $\text{Div}_+^g(\Gamma) \rightarrow \text{Jac}(\Gamma)$ must consist of break divisors outside the $(g - 1)$ -skeleton of the ABKS decomposition. The rest of the section will be uniquely determined by taking limit from the generic part. \square

Now we explain how the crucial observation in the proof of Proposition 3.3.2, namely that most fibers of $\text{Div}_+^g(\Gamma) \rightarrow \text{Jac}(\Gamma)$ are singletons, can be viewed as an instance of a

tropical analogue of Sard's theorem (we will be referring to basic terminology in differential geometry, which can be found in standard texts such as [81]).

Since $\text{Div}_+^g(\Gamma)$ has no canonical tropical structure, we will work with $\Gamma^g \rightarrow \text{Jac}(\Gamma)$ instead. Fix a model G of Γ , denote the edge set of G by E . For a point $x \in \Gamma$, the *tangent space* $T_x\Gamma \subset \mathbb{R}^E$ is the union of cones $\text{Cone}(\vec{e})$ corresponding to the edge directions \vec{e} 's emanating from x . For example, if x is in the interior of an edge e , then $T_x\Gamma$ is the line spanned by e as the union of two rays, each corresponds to a direction of e emanating from x .

Next we define the tangent space $T_{(p_1, \dots, p_g)}\Gamma^g$ at a point (p_1, \dots, p_g) of Γ^g , which will be a fan in $\mathbb{R}^{\{1, \dots, g\} \times E}$. For each $i = 1, \dots, g$, we identify $T_{p_i}\Gamma$ as a fan in the i -th copy of \mathbb{R}^E , then we take the Minkowski sum of the fans $T_{p_i}\Gamma$'s. Equivalently, $T_{(p_1, \dots, p_g)}\Gamma^g$ consists of cones (and faces thereof) of the form $\text{Cone}(\{1\} \times \vec{e}_1, \dots, \{g\} \times \vec{e}_g)$, where each \vec{e}_i is an edge direction emanating from p_i .

Remark. Our definition is similar to the case in differential geometry, where each point x in a manifold X admits a neighborhood diffeomorphic to a neighborhood of its tangent space. For the sake of simplicity, we do not consider the multiplicity data on Γ^g or $T_{(p_1, \dots, p_g)}\Gamma^g$.

Lastly, since $\text{Jac}(\Gamma) = H_1(\Gamma; \mathbb{R})/H_1(\Gamma; \mathbb{Z})$, we define the tangent space $T_{[D]}\text{Jac}(\Gamma)$ at any point of $\text{Jac}(\Gamma)$ as the cycle space $H_1(\Gamma; \mathbb{R})$ itself. Now we can define the *derivative* of the $\Gamma^g \rightarrow \text{Jac}(\Gamma)$ map as the restriction (to $T_{(p_1, \dots, p_g)}\Gamma^g$) of the linear map $\mathbb{R}^{\{1, \dots, g\} \times E} \rightarrow H_1(\Gamma; \mathbb{R})$, which extends the orthogonal projection $\mathbb{R}^E \cong C_1(\Gamma; \mathbb{R}) \rightarrow H_1(\Gamma; \mathbb{R})$ of each \mathbb{R}^E -summand linearly.

Proposition 3.3.3 *An effective degree g divisor $(p_1) + \dots + (p_g)$ is break if and only if the image of $T_{(p_1, \dots, p_g)}\Gamma^g$ under the derivative map has a full dimensional image.*

PROOF: By definition, the full-dimensional cones in $T_{(p_1, \dots, p_g)}\Gamma^g$ (and the extremal rays thereof) specify all possible ways to choose a multi-set of g directed edges from G whose

“orientable divisor” is $(p_1) + \dots + (p_g)$. The projection of g edges span $H_1(\Gamma; \mathbb{R})$ if and only if the edges are the complement of some spanning tree. Therefore $(p_1) + \dots + (p_g)$ is break if and only if some cone of $T_{(p_1, \dots, p_g)}\Gamma^g$ has a full-dimensional image under the derivative map. \square

Notice that the proof of Proposition 3.3.3 is similar to the proof of Proposition 3.2.1: $T_{(p_1, \dots, p_g)}\Gamma^g$ encodes information of the transversal matroid while $T_{[D]}\text{Jac}(\Gamma)$ encodes information of the cographic matroid, and having full-dimensional image implies the two matroids are having common bases. It is an interesting problem to further investigate the relations between geometry of tropical objects and matroidal operations.

Further mimicking notions in differential geometry, we can say that a point (p_1, \dots, p_g) in Γ^g is a *critical point* of f whenever $(p_1) + \dots + (p_g)$ is not break, and a degree g divisor class in $\text{Jac}(\Gamma)$ is a *critical value* if it contains a non-break effective divisor. The fact that the subset of critical values is contained in a codimension one subset (the $(g - 1)$ -skeleton of the ABKS decomposition) can be viewed as an analogue of Sard’s theorem [81, Theorem 10.7].

3.4 The f -Vector of the ABKS Decomposition

We study the f -vector of the ABKS decomposition of a graph G , that is, the number of faces of different dimensions. Denote by Γ the metric realization of G whose all edges lengths are 1.

Definition 3.4.1 *A pair (D', E') is a break (i, j) -configuration if $D' \in \text{Div}^i(G)$ and $E' \subset E(G)$ is a set of j edges such that $G - E'$ is connected and D' , treated as a divisor on $G \setminus E'$, is break. It follows that we must have $i + j = g$.*

Each break divisor D on Γ corresponds to a break configuration (D', E') , where D' is D restricted on $V(G)$, and E' is the set of edges containing chips of $D - D'$ in their

interior. Conversely every break configuration $(D', E' = \{e_1, \dots, e_j\})$ corresponds to the family $\{D' + (v_1) + \dots + (v_j) : v_i \in e_i, \forall i = 1, 2, \dots, j\}$ of break divisors.

Proposition 3.4.2 *For each $0 \leq i \leq g$, there is a one-to-one correspondence between the i -dimensional faces of the ABKS decomposition and the break $(g - i, i)$ -configurations.*

PROOF: The assertion follows from unwrapping the construction of the ABKS decomposition described in Section 3.1.3. Every face of the ABKS decomposition is the face of some cell $\mathcal{C}_T \cong f_1 \times \dots \times f_g$, where f_1, \dots, f_g are the edges not in the spanning tree T . An i -dimensional face of \mathcal{C}_T is specified by choosing endpoints for $g - i$ edges from f_i 's, which specifies a break $(g - i, i)$ -configuration. Conversely, given a break $(g - i, i)$ -configuration (D', E') , pick a spanning tree $T \subset G - E'$, then some i -dimensional face of the cell \mathcal{C}_T corresponds to such configuration. \square

Theorem 3.4.3 *The f -vector of the ABKS decomposition of G is $|\mathcal{ST}(G)| \cdot \binom{g}{0}, \binom{g}{1}, \dots, \binom{g}{g}$. Recall that $|\mathcal{ST}(G)|$ is the number of spanning trees of G .*

PROOF: By Proposition 3.4.2 it suffices to count the number of break $(g - i, i)$ -configurations. Fix a bijection φ_H between $\mathcal{BD}(H)$ and $\mathcal{ST}(H)$ for every connected spanning subgraph H of G . We associate each break $(g - i, i)$ -configuration (D', E') with a pair $(T = \varphi_{G-E'}(D'), E')$; conversely, for every pair (T, E') , where $|E'| = i$ and T is a spanning tree of $G - E'$ (hence G), we can associate a break $(g - i, i)$ -configuration $(\varphi_{G-E'}^{-1}(T), E')$. Since the two maps are inverse of each other, we conclude that the number of break $(g - i, i)$ -configurations is the same as the number of these (T, E') pairs, and there are exactly $\binom{g}{i} |\mathcal{ST}(G)|$ of them: for each spanning tree T of G , any i -subset E' of $E(G) \setminus E(T)$ could pair up with T . \square

Remark. The combinatorics of break configurations has connections with classical algebraic geometry, notably the theory of *compactified Jacobians*. Let X be a nodal curve. Let G be the *dual graph* of X , where vertices are the irreducible components of X , and

two vertices are adjacent if the corresponding components intersect. The *Jacobian* $\text{Jac}(X)$ of X is the moduli space of isomorphism classes of degree zero line bundles on X , and one can compactify $\text{Jac}(X)$ by including suitable generalization of line bundles to it. The *tropicalization* of these algebro-geometric objects can be recorded as divisors on G (or metric realization thereof). In the case of Simpson compactifications [100], the combinatorial type of such divisors are the break configurations, and as a corollary, the compactifications are independent of some initial choices [99]. On the other hand, Karl Christ independently computed the f -vector of the ABKS decomposition, and used it to compute topological invariants of the Simpson compactification and related objects [29].

Finally, we give one example of how the combinatorics of break divisors and the ABKS decomposition could be used to study tropical Brill-Noether theory. Similar to the case of Proposition 3.3.1, the following statement is related to [88, Theorem 8.5], but with a more transparent proof.

Proposition 3.4.4 *The Brill-Noether locus W_g^1 is contained in the $(g - 2)$ -skeleton of the ABKS decomposition. In particular, $\dim(W_g^1) \leq \rho(g, g, 1) = g - 2$.*

PROOF: Let D be a break divisor outside the $(g - 2)$ -skeleton, we claim it is of rank 0 by showing that D is q -reduced for some $q \in \Gamma$ that has no chips. By Proposition 3.4.2, there are at least $g - 1$ chips located in the interior of edges. Choose any $g - 1$ of them and remove the corresponding edges from G , there is a unique cycle C left in the graph, and the remaining chip x is on C . Pick any point q on $C \setminus \{x\}$. Apply Dhar's burning algorithm and start a fire at a point q , the fire can burn through the whole of C because it will approach x from the two sides, from there the rest of the graph will be burnt using a similar argument as in Lemma 3.3.1. □

3.5 A Recursive Proof of the Number of (Generalized) Break Divisors

We consider a generalization of (the set of) break divisors on graphs and show that the cardinality of any such set of divisors is at least the number of spanning trees of a graph. We further show that the inequality is an equality when we consider ordinary break divisors. This gives a new, recursive proof for Part (4) of Proposition 3.1.2 without referring to [1] or the Matrix–Tree theorem.

The proof presented here was first given by myself for equality of the ordinary break divisors case. Jonathan Barmak later independently used an argument equivalent to the proof of Proposition 3.5.3 to prove the inequality for the general case.

Recall that a break divisor $D \in \mathcal{BD}(G)$ is induced by a spanning tree T and an orientation \mathcal{O} of edges outside T , denoted as $D = \mathcal{D}(T, \mathcal{O})$, if D is the indegree sequence of the partial orientation \mathcal{O} . Let I_G be an arbitrary function from the $\mathcal{ST}(G)$ to $\text{Div}^0(G)$. A divisor of degree g is a *generalized break divisor with respect to I_G* if it can be written as $\mathcal{D}(T, \mathcal{O}) + I_G(T)$ for some choice of spanning tree $T \in \mathcal{ST}(G)$ and orientation \mathcal{O} of edges outside T . Denote by $\mathcal{BD}_{I_G}(G)$ the set of such divisors.

Theorem 3.5.1 *For any function I_G , $|\mathcal{BD}_{I_G}(G)| \geq |\mathcal{ST}(G)|$. Equality holds when $I_G \equiv 0$.*

Theorem 3.5.1 was motivated by a question posed by Jesse Kass and Nicola Pagani in their work on compactified Jacobians [74]. In their construction of compactified Jacobians of a nodal curve X with dual graph G , one chooses a set of $|\mathcal{ST}(G \setminus F)|$ divisors from $\text{Div}^{g-|F|}(G)$ for every subset F of $E(G)$. Since one only needs to consider subsets in which $G \setminus F$ are connected, the function I_G records the choices for those maximal subsets. The divisors chosen correspond to the strata of certain stratification of the moduli space; and in order the moduli space to be well-defined, whenever a divisor D is chosen for F , the divisor D' obtained by adding a chip to D at an endpoint of some $f \in F$ must also be chosen for $F \setminus \{f\}$. The inequality in Theorem 3.5.1 roughly says that one has enough combinatorial options for choosing divisors for $F = \emptyset$, starting with any I_G .

To establish the recursive proof, we fix an edge $e = uv$ of G that is neither a loop nor an isthmus; the case when G consists solely of loops and isthmuses will be handled separately as a base case. For $D \in \text{Div}(G)$, denote by $f(D) \in \text{Div}(G/e)$ the divisor such that $f(D)(x) = D(x)$ for $x \in V(G) \setminus \{u, v\}$ and $f(D)(w) = D(u) + D(v)$, where w corresponds to the contracted edge e . Now define $I_{G \setminus e}(T) := I_G(T)$ and $I_{G/e}(T) := f(I_G(T \cup e))$.

Lemma 3.5.2 *If $E \in \text{Div}^{g-1}(G \setminus e)$ is a generalized break divisor with respect to $I_{G \setminus e}$, then both $E + (u)$ and $E + (v)$ are generalized break divisors with respect to I_G .*

PROOF: Suppose $E = \mathcal{D}(T, \mathcal{O}) + I_{G \setminus e}(T) \in F(G \setminus e)$ for spanning tree $T \in \mathcal{ST}(G \setminus e)$ and orientation \mathcal{O} of edges (of $G \setminus e$) outside T . Then $E + (u) = \mathcal{D}(T, \mathcal{O} \cup \vec{vu}) + I_G(T) \in \mathcal{BD}_{I_G}(G)$ and $E + (v) = \mathcal{D}(T, \mathcal{O} \cup \vec{uv}) + I_G(T) \in \mathcal{BD}_{I_G}(G)$. \square

Define the function $\gamma : \mathcal{BD}_{I_G}(G) \rightarrow \mathcal{BD}_{I_{G \setminus e}}(G \setminus e) \sqcup \mathcal{BD}_{I_{G/e}}(G/e)$ as follows:

1. If $D - (v)$ (considered as a divisor on $G \setminus e$) is in $\mathcal{BD}_{I_{G \setminus e}}(G \setminus e)$, then set $\gamma(D) := D - (v) \in \mathcal{BD}_{I_{G \setminus e}}(G \setminus e)$;
2. else if $f(D) \in \mathcal{BD}_{I_{G/e}}(G/e)$, then set $\gamma(D) := f(D) \in \mathcal{BD}_{I_{G/e}}(G/e)$;
3. otherwise set $\gamma(D) := D - (u) \in \mathcal{BD}_{I_{G \setminus e}}(G \setminus e)$.

Proposition 3.5.3 *The function γ is well-defined and surjective.*

PROOF: “ γ is well-defined”. Suppose $D \in \mathcal{BD}_{I_G}(G)$ is equal to $\mathcal{D}(T, \mathcal{O}) + I_G(T)$. Then either

1. $e \notin T$ and e is oriented as \vec{uv} in \mathcal{O} , which means $D - (v) = \mathcal{D}(T, \mathcal{O} \setminus \vec{uv}) + I_{G \setminus e}(T) \in \mathcal{BD}_{I_{G \setminus e}}(G \setminus e)$, or
2. $e \in T$, which means $f(D) = \mathcal{D}(T \setminus e, \mathcal{O}) + I_{G/e}(T \setminus e) \in \mathcal{BD}_{I_{G/e}}(G/e)$, or

3. $e \notin T$ and e is oriented as \vec{vu} in \mathcal{O} , which means $D-(u) = \mathcal{D}(T, \mathcal{O} \setminus \vec{vu}) + I_{G \setminus e}(T) \in \mathcal{BD}_{I_{G \setminus e}}(G \setminus e)$.

So $\gamma(D)$ always lands in $\mathcal{BD}_{I_{G \setminus e}}(G \setminus e) \sqcup \mathcal{BD}_{I_{G/e}}(G/e)$.

“ γ is surjective on $\mathcal{BD}_{I_{G \setminus e}}(G \setminus e)$ ”. For each $D \in \mathcal{BD}_{I_{G \setminus e}}(G \setminus e)$, $D+(v) \in \mathcal{BD}_{I_G}(G)$ by Lemma 3.5.2, therefore $\gamma(D+(v)) = D$ as we are in the first case in the definition of γ .

“ γ is surjective on $\mathcal{BD}_{I_{G/e}}(G/e)$ ”. For each $\tilde{D} = \mathcal{D}(\tilde{T}, \tilde{\mathcal{O}}) + I_{G/e}(\tilde{T}) \in \mathcal{BD}_{I_{G/e}}(G/e)$, set $D := \mathcal{D}(\tilde{T} \cup e, \tilde{\mathcal{O}}) + I_G(\tilde{T} \cup e) \in \mathcal{BD}_{I_G}(G)$ and let K be the minimum $k \geq 0$ such that $D-(k+1)(v)+(k+1)(u)$ is not in $\mathcal{BD}_{I_G}(G)$. Write $D' := D-K(v)+K(u) \in \mathcal{BD}_{I_G}(G)$. If $D'-(v) \in \mathcal{BD}_{I_{G \setminus e}}(G \setminus e)$, then by Lemma 3.5.2, $D'-(v)+(u) \in \mathcal{BD}_{I_G}(G)$, a contradiction to the choice of K . Therefore we are not in the first case, and $\gamma(D') = f(D') = f(D) = \tilde{D}$. \square

Now we show that γ is injective as well when I_G is identically the zero divisor, that is, when $\mathcal{BD}_{I_G}(G)$ is the set of ordinary break divisors.

Proposition 3.5.4 *The function γ is injective when $I_G \equiv 0$.*

PROOF: “ γ is injective in the first case”. Obvious.

“ γ is injective in the second case”. Let $D \in \mathcal{BD}_0(G)$ be a break divisor such that $\gamma(D) \in \mathcal{BD}_0(G/e)$. Since $D-(v)$ is not a break divisor on $G \setminus e$ (or otherwise we would have been in the first case), there exists a non-empty subset $S_0 \subset V(G)$ such that $\chi_{D-(v), G \setminus e}(S_0) \leq 0$ by Lemma 3.1.3. Note that $\chi_{D-(v), G \setminus e}(S) = \chi_{D, G}(S) > 0$ whenever $S \cap \{u, v\} \neq \{v\}$, so we must have $S_0 \cap \{u, v\} = \{v\}$ and $\chi_{D, G}(S_0) = 1$. From this we can see that for any $k > 0$, $D-k(v)+k(u)$ is not a break divisor on G , for $\chi_{D-k(v)+k(u), G}(S_0) \leq \chi_{D, G}(S_0) - 1 \leq 0$. Reversing the roles in the argument, we can also see that $D+k(v)-k(u)$'s are either not break (because it fails the criterion in Lemma 3.1.3) or would not be mapped to $\mathcal{BD}(G/e)$ (because $D-(v)$ is already break by Lemma 3.1.3). But from the

definition of γ , these are the only possible divisors that could be mapped to D via f , so γ is injective on the second part.

“ γ is injective in the third case”. We claim that such case can never happen, hence injectivity is vacuously true. Suppose $D \in \mathcal{BD}_0(G)$ is equal to $\mathcal{D}(T, \mathcal{O})$ with $e \notin T$ and e is oriented as \vec{vu} in \mathcal{O} . Pick, from the fundamental cycle of e with respect to T , the other edge f incident to u and set $T' = T - f + e$. Construct the partial orientation \mathcal{O}' of G that agrees with \mathcal{O} everywhere except e is now unoriented and f is oriented towards u . Then $D = (T', \mathcal{O}')$, and we would have stopped by the latest the second case. \square

Remark. Despite not using the terminology of Euler characteristic, several other bijective proofs related to break divisors [13, 114] construct such S_0 explicitly in their arguments.

PROOF OF THEOREM 3.5.1: If a graph only consists of loops and isthmuses, then it only has one spanning tree and one break divisor, and the theorem is true with equality. By Proposition 3.5.3, for a graph with an edge e that is neither a loop nor an isthmus, $|\mathcal{BD}_{I_G}(G)| \geq |\mathcal{BD}_{I_{G \setminus e}}(G \setminus e)| + |\mathcal{BD}_{I_{G/e}}(G/e)|$, and by induction, $|\mathcal{BD}_{I_{G \setminus e}}(G \setminus e)| + |\mathcal{BD}_{I_{G/e}}(G/e)| \geq |\mathcal{ST}(G \setminus e)| + |\mathcal{ST}(G/e)| = |\mathcal{ST}(G)|$, with equality through out if $I_G \equiv 0$ by Proposition 3.5.4. \square

CHAPTER 4

CIRCUIT-COCIRCUIT REVERSAL SYSTEMS

4.1 Definition

Let M be an oriented matroid on E . An orientation \mathcal{O} is *compatible* with a signed circuit or signed cocircuit C if $\mathcal{O}(e) = C(e)$ for every $e \in \underline{C}$. We have the following dichotomy.

Proposition 4.1.1 [23, Corollary 3.4.6] *In an orientation \mathcal{O} , every element is either contained in a signed circuit compatible with \mathcal{O} or a signed cocircuit compatible with \mathcal{O} , but not both.*

The set of elements contained in some signed circuit (resp. signed cocircuit) is the *circuit part* (resp. *cocircuit part*) of \mathcal{O} . An orientation is *acyclic* if its cocircuit part is the all of E , and it is *totally cyclic* if its circuit part is the all of E .

The following definitions were introduced by Gioan [57, 58].

Definition 4.1.2 *Given an orientation \mathcal{O} , a circuit reversal (resp. cocircuit reversal) picks a signed circuit (resp. cocircuit) C compatible with \mathcal{O} and reverses the sign of $\mathcal{O}(e)$ for every $e \in \underline{C}$.*

Two orientations $\mathcal{O}_1, \mathcal{O}_2$ are circuit-cocircuit reversal equivalent if \mathcal{O}_1 can be transformed into \mathcal{O}_2 via a sequence of circuit and cocircuit reversals. The set of equivalence classes $\mathcal{G}(M)$ is the circuit-cocircuit reversal system of M , and each equivalence class is a circuit-cocircuit reversal class. Similarly we can define the circuit reversal classes or cocircuit reversal classes of M .

Lemma 4.1.3 [58, Proposition 1] *A circuit (resp. cocircuit) reversal does not change the circuit part nor the cocircuit part of an orientation.*

Gioan proved the following theorem, which can be thought as a variation of the Matrix–Tree theorem.

Theorem 4.1.4 [58, Theorem 10(v)] *Let M be a regular matroid. Then $|\mathcal{G}(M)| = |\mathcal{B}(M)|$.*

Built on the work of Gioan, Backman established the following connection between circuit-cocircuit reversals and chip-firing in the case of graphs.

Proposition 4.1.5 [4, Lemma 3.1, Theorem 3.3], [57, Proposition 4.10, Corollary 4.13] *Let $\mathcal{O}, \mathcal{O}'$ be two orientations of a graph G . Then $D_{\mathcal{O}} = D_{\mathcal{O}'}$ if and only if \mathcal{O}' can be obtained from \mathcal{O} by a sequence of circuit (cycle) reversals, and $D_{\mathcal{O}} \sim D_{\mathcal{O}'}$ if and only if \mathcal{O}' can be obtained from \mathcal{O} by a sequence of circuit (cycle) reversals and cocircuit (cut) reversals.*

Proposition 4.1.6 [4, Section 5] *The map $\tau_G : \mathcal{G}(M(G)) \rightarrow \text{Pic}^{g-1}(G)$ given by $[\mathcal{O}] \mapsto [D_{\mathcal{O}}]$ is a well-defined bijection.*

For the rest of Section 4, we will use the term *positive* (co)circuit (with respect to an orientation \mathcal{O}) to denote either a signed (co)circuit that is compatible with \mathcal{O} , or a (co)circuit C of \underline{M} in which $\mathcal{O}|_C$ is a signed (co)circuit in M . Furthermore, given an orientation \mathcal{O} and a subset $X \subset E$, denote by ${}_{-X}\mathcal{O}$ the orientation obtained from reversing elements of X in \mathcal{O} . For a (co)circuit C of \underline{M} , we say that ${}_{-X}C$ is positive with respect to \mathcal{O} if C is a positive (co)circuit of ${}_{-X}\mathcal{O}$. Finally, denote by χ_X the $\{0, 1\}$ -vector whose support is X .

4.2 Circuit-cocircuit Reversal System as a $\text{Jac}(M)$ -Torsor

We will define a natural action of $\text{Jac}(M)$ on the circuit-cocircuit reversal system $\mathcal{G}(M)$ for a regular matroid M , and prove that the action is simply transitive. We will also give an efficient algorithm for computing this action. For the rest of Section 4.2, we will fix a full row rank totally unimodular matrix A realizing M .

4.2.1 Description of the Action

Recall from Definition 2.4.1 that $\text{Jac}(M)$ can be identified with $\frac{C_1(M)}{Z_1(M) \oplus B_1(M)}$. The quotient group is generated by $[\vec{e}]$, $e \in M$, where we use an overhead arrow to emphasize that we are keeping track of the orientation of elements.

The group action $\text{Jac}(M) \circ \mathcal{G}(M)$ is defined by linearly extending the following action of each generator $[\vec{e}]$ on circuit-cocircuit reversal classes: given a class, pick an orientation \mathcal{O} from the class so that e is oriented as \vec{e} in \mathcal{O} , reverse the orientation of e in \mathcal{O} to obtain \mathcal{O}' , and set $[\vec{e}] \cdot [\mathcal{O}] = [\mathcal{O}']$. This action generalizes the one defined in terms of *path reversals* by Backman in the graphical case [4, Section 5].

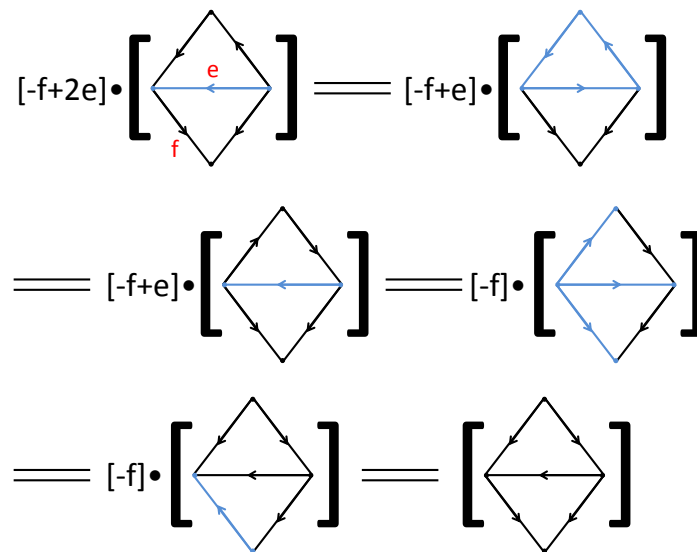


Figure 4.1: Example of the torsor. Here the reference orientations of e, f are the same as the orientation we begin with.

We will prove the following theorem.

Theorem 4.2.1 *The group action \circ is well-defined and simply transitive.*

Theorem 4.2.1 will be deduced from a series of intermediate results. We start with a few general lemmas regarding (regular) matroids.

Lemma 4.2.2 *Let $e \in E$, and suppose $X \subset E \setminus \{e\}$ is a positive cocircuit of $\mathcal{O} \setminus e$ but not in \mathcal{O} . Then $Y := X \cup \{e\}$ is a positive cocircuit of either \mathcal{O} or ${}_{-e}\mathcal{O}$.*

PROOF: Without loss of generality we may assume $\mathcal{O} \equiv +$, hence there exists some $\mathbf{u} \in \mathbb{Z}^r$ such that $\mathbf{u}^T A|_{E \setminus \{e\}} = \chi_X$. Now $\mathbf{u}^T A = \chi_X + \lambda \chi_{\{e\}}$ for some λ , which is not zero as X is not a cocircuit in M . Since the support of $\mathbf{u}^T A$ is Y , Y contains a cocircuit D by the dual of Lemma 2.3.9.

If $D \cap X = \emptyset$, then $D = \{e\}$. So $\mathbf{u}^T A = \chi_{\{e\}}$ for some \mathbf{u}' , which implies that $(\mathbf{u} - \lambda \mathbf{u}')^T A = \chi_X$, thus X itself is a cocircuit in M , a contradiction. Now we must have $X \subset D$, or otherwise $D \cap X \subsetneq X$ would be a cocircuit in $M \setminus e$. Since $D \neq X$, $Y = D$ is a cocircuit, and depending the sign of λ , either Y or ${}_{-e}Y$ is positive in \mathcal{O} . \square

Lemma 4.2.3 *Suppose $e \in E$ is contained in some positive circuit of \mathcal{O} , and Y is a subset of E containing e such that ${}_{-e}Y$ is a positive cocircuit of \mathcal{O} . Then any positive circuit containing e intersects Y in exactly two elements.*

PROOF: Again we assume $\mathcal{O} \equiv +$. Let C be a positive circuit containing e , so $A\chi_C = 0$. By assumption, there exists a vector \mathbf{u} such that $\mathbf{u}^T A = \chi_{Y \setminus \{e\}} - \chi_{\{e\}}$. Then $0 = \mathbf{u}^T A\chi_C = |(Y \setminus \{e\}) \cap C| - 1$, that is, Y intersects C in e together with exactly one more element. \square

4.2.2 The Action is Well-defined

In order to show that the action of $\text{Jac}(M)$ on $\mathcal{G}(M)$ is well-defined, we first show that the corresponding action (which by abuse of notation we continue to write as \circlearrowleft) of $C_1(M)$ on $\mathcal{G}(M)$ is well-defined, then that the action descends to the quotient by $Z_1(M) \oplus B_1(M)$.

Proposition 4.2.4 *For every $[\mathcal{O}] \in \mathcal{G}(M)$ and oriented element \vec{e} , there exists $\tilde{\mathcal{O}} \in [\mathcal{O}]$ so that e is oriented as \vec{e} in $\tilde{\mathcal{O}}$.*

PROOF: This follows from Proposition 4.1.1, which guarantees that e is either contained in a positive circuit or cocircuit C of \mathcal{O} . If e is not already oriented as \vec{e} in \mathcal{O} , reverse C . \square

Proposition 4.2.5 *The action of \vec{e} on $[\mathcal{O}]$ is independent of which orientation we choose.*

PROOF: It suffices to show that if $\mathcal{O} \sim \mathcal{O}'$ and the two orientations agree on e , then ${}_{-e}\mathcal{O} \sim {}_{-e}\mathcal{O}'$. By Lemma 2.3.14 and its dual, \mathcal{O} and \mathcal{O}' differ by a disjoint union of positive circuits and cocircuits which do not contain e , and ${}_{-e}\mathcal{O}$ can be transformed to ${}_{-e}\mathcal{O}'$ by reversing these circuits and cocircuits. \square

Proposition 4.2.6 *For any $\vec{e}, \vec{f} \in C_1(M)$ and $[\mathcal{O}] \in \mathcal{G}(M)$, $\vec{e} \cdot (\vec{f} \cdot [\mathcal{O}]) = \vec{f} \cdot (\vec{e} \cdot [\mathcal{O}])$. Hence it is valid to extend \cdot linearly, and \circlearrowleft is indeed a group action of $C_1(M)$ on $\mathcal{G}(M)$.*

PROOF: The statement is tautological if $\vec{e} = \vec{f}$. If $\vec{e} = -\vec{f}$, then without loss of generality the orientation of e in \mathcal{O} is \vec{e} . Let C be a positive (co)circuit containing e . On one hand, we have $\vec{f} \cdot (\vec{e} \cdot [\mathcal{O}]) = \vec{f} \cdot [{}_{-e}\mathcal{O}] = [\mathcal{O}]$; on the other hand, $\vec{e} \cdot (\vec{f} \cdot [\mathcal{O}]) = \vec{e} \cdot (\vec{f} \cdot [{}_{-C}\mathcal{O}]) = [{}_{-C}\mathcal{O}] = [\mathcal{O}]$. Therefore the action order does not matter.

Now suppose $e \neq f$. We may again assume without loss of generality that e is oriented as \vec{e} in \mathcal{O} . The statement is easy if there exists some positive (co)circuit in \mathcal{O} that contains f but not e , as we can reverse it and obtain an orientation in which the orientations of e, f are already \vec{e}, \vec{f} . So without loss of generality e, f are in the circuit part of \mathcal{O} and every positive circuit containing f also contains e . Fix any such positive circuit C . f must be in some positive cocircuit D' of $\mathcal{O} \setminus e$, or otherwise f is in some positive circuit of $\mathcal{O} \setminus e$, which is a positive circuit of \mathcal{O} avoiding e . By Lemma 4.2.2, $D := D' \cup \{e\}$ is a cocircuit in \mathcal{O} and ${}_{-e}D$ is positive, and by Lemma 4.2.3, we know that $C \cap D = \{e, f\}$.

We have $\vec{e} \cdot [\mathcal{O}] = [{}_{-e}\mathcal{O}] = [{}_{-(D \setminus \{e\})}\mathcal{O}]$ as D is a positive cocircuit of ${}_{-e}\mathcal{O}$, and $\vec{f} \cdot [{}_{-(D \setminus \{e\})}\mathcal{O}] = [{}_{-(D \setminus \{e, f\})}\mathcal{O}]$ as f is oriented as \vec{f} in ${}_{-(D \setminus \{e\})}\mathcal{O}$. On the other hand, $\vec{f} \cdot [\mathcal{O}] = \vec{f} \cdot [{}_{-C}\mathcal{O}] = [{}_{-(C \setminus \{f\})}\mathcal{O}]$, and D is positive in ${}_{-(C \setminus \{f\})}\mathcal{O}$ since $C \cap D = \{e, f\}$, hence $\vec{e} \cdot [{}_{-(C \setminus \{f\})}\mathcal{O}] = \vec{e} \cdot [{}_{-(C \cup D \setminus \{e\})}\mathcal{O}] = [{}_{-(C \cup D)}\mathcal{O}]$. But C is positive in ${}_{-(C \cup D)}\mathcal{O}$, so $[{}_{-(C \cup D)}\mathcal{O}] = [{}_{-(C \cup D) \Delta C}\mathcal{O}] = [{}_{-(D \setminus \{e, f\})}\mathcal{O}]$. \square

Now we know that $C_1(M) \circlearrowleft \mathcal{G}(M)$ is well-defined, so we show next that this action descends to a group action $\text{Jac}(M) \circlearrowleft \mathcal{G}(M)$.

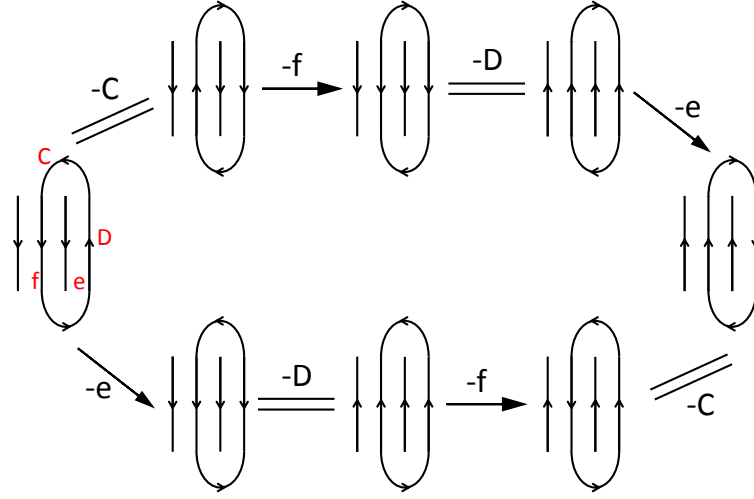


Figure 4.2: Illustration of Proposition 4.2.6.

Proposition 4.2.7 *The stabilizer of the action on any $[\mathcal{O}]$ contains $Z_1(M) \oplus B_1(M)$.*

PROOF: Let $\vec{C} \in Z_1(M)$ be a signed circuit. Let F be the set of elements in C whose orientation in \mathcal{O} is the same as in \vec{C} . Then $\vec{C} \cdot [\mathcal{O}] = (\sum_{\vec{e} \in \vec{C} \setminus F} \vec{e}) \cdot [-F\mathcal{O}] = (\sum_{\vec{e} \in \vec{C} \setminus F} \vec{e}) \cdot [-_{(C \setminus F)}\mathcal{O}] = [\mathcal{O}]$. The proof for $B_1(M)$ is similar. \square

4.2.3 The Action is Simply Transitive

Proposition 4.2.8 *The group action $\text{Jac}(M) \curvearrowright \mathcal{G}(M)$ is transitive.*

PROOF: Given any two orientations $\mathcal{O}, \mathcal{O}'$, let γ be the sum of the oriented elements in \mathcal{O} whose orientation in \mathcal{O}' is different; then $[\gamma] \cdot [\mathcal{O}] = [\mathcal{O}']$. \square

PROOF OF THEOREM 4.2.1: We know from Proposition 4.2.7 that $\text{Jac}(M) \curvearrowright \mathcal{G}(M)$ is a well-defined group action, and which by Proposition 4.2.8 is transitive. Since $|\text{Jac}(M)| = |\mathcal{B}(M)| = |\mathcal{G}(M)|$, the action is automatically simple. \square

It is worthwhile to give a direct proof of the simplicity of the action which does not make use of the equality $|\text{Jac}(M)| = |\mathcal{G}(M)|$, since this yields an independent and “bijective” proof of the equality. We begin with the following reduction.

Proposition 4.2.9 *The simplicity of the group action $\text{Jac}(M) \curvearrowright \mathcal{G}(M)$ is equivalent to the statement that every element of the quotient group $\frac{C_1(M)}{Z_1(M) \oplus B_1(M)}$ contains a coset representative whose coefficients are all 1, 0, -1 .*

PROOF: Suppose the group action is simple. Let $[\gamma] \in \frac{C_1(M)}{Z_1(M) \oplus B_1(M)}$ be an element of $\text{Jac}(M)$. Pick an arbitrary orientation \mathcal{O} and an arbitrary orientation \mathcal{O}' from $[\gamma] \cdot [\mathcal{O}]$. Let $\gamma_0 \in C_1(M)$ be the sum of oriented elements in \mathcal{O} which have the opposite orientation in \mathcal{O}' . Then $[\gamma_0] \cdot [\mathcal{O}] = [\mathcal{O}']$, which by the simplicity of the action implies that $\gamma_0 \in [\gamma]$. The desired coset representative is γ_0 .

Conversely, suppose such a set of coset representatives exists. We need to show that whenever $[\gamma] \in \frac{C_1(M)}{Z_1(M) \oplus B_1(M)}$ fixes some circuit-cocircuit reversal class $[\mathcal{O}]$, $[\gamma] = [0]$. By transitivity, we may assume $[\gamma]$ fixes all circuit-cocircuit reversal classes: for any other reversal class $[\mathcal{O}'] = [\gamma'] \cdot [\mathcal{O}]$, $[\gamma] \cdot [\mathcal{O}'] = ([\gamma'] + [\gamma] + [-\gamma']) \cdot [\mathcal{O}'] = ([\gamma'] + [\gamma]) \cdot [\mathcal{O}'] = [\gamma'] \cdot [\mathcal{O}'] = [\mathcal{O}']$. Without loss of generality, the coefficients of γ are all 1, 0, -1 with support $F \subset E$. Pick an orientation \mathcal{O} in which the orientation of every element in F agrees with γ , then $[\mathcal{O}] = [\gamma] \cdot [\mathcal{O}] = [-_F \mathcal{O}]$. Therefore $\mathcal{O} \sim_{-F} \mathcal{O}$, meaning that F is a disjoint union of positive circuits and cocircuits in \mathcal{O} , that is, $\gamma \in Z_1(M) \oplus B_1(M)$ and $[\gamma] = [0]$. \square

Then we prove the claim about special representatives in $\text{Jac}(M)$.

Proposition 4.2.10 *Every element of $\frac{C_1(M)}{Z_1(M) \oplus B_1(M)}$ contains a coset representative whose coefficients are all 1, 0, -1 .*

PROOF: We will show that there is such a representative in $[\gamma]$ for every $\gamma = \sum_{e \in E} c_e e \in C_1(M)$ by lexicographic induction on $|\gamma|_\infty := \max_{e \in E} |c_e|$ and the number of elements e with $|c_e| = |\gamma|_\infty$. The assertion is clearly true if $|\gamma|_\infty \leq 1$, so suppose $|\gamma|_\infty > 1$.

By reorientation, we may assume that all coefficients of γ are non-negative. Pick an element e whose coefficient c_e equals $|\gamma|_\infty$. By applying Proposition 4.1.1 to the all positive orientation, there exists a signed (co)circuit $C \geq 0$ containing e . Now if we subtract $\gamma_C := \sum_{f \in C} f$ from γ , all positive coefficients c_f with $f \in C$ decrease by 1, while the zero coefficients in the support of C become -1 . Hence $|\gamma - \gamma_C|_\infty \leq |\gamma|_\infty$ and the number of elements f with $|c_f| = |\gamma|_\infty$ strictly decreases. By our induction hypothesis, there exists a representative of the desired form in $[\gamma - \gamma_C] = [\gamma]$. \square

Remark. Proposition 4.2.10 implies the following result for divisors on graphs, which serves as a degree 0 analogue of Proposition 4.1.6: for every degree 0 divisor D , there exists a partial orientation \mathcal{O} of G such that $D \sim D'$, where $D'(v) = \text{indeg}_{\mathcal{O}}(v) - \text{outdeg}_{\mathcal{O}}(v), \forall v \in V(G)$.

4.2.4 The Action is Efficiently Computable

Finally, we show that the simply transitive action of $\text{Jac}(M)$ on $\mathcal{G}(M)$ is efficiently computable.

Proposition 4.2.11 *The action of $\text{Jac}(M)$ on $\mathcal{G}(M)$ can be computed in polynomial-time, given a totally unimodular matrix A realizing M .*

PROOF: First we note that computing the action of a generator $[\vec{e}]$ on a circuit-cocircuit reversal class can be done in polynomial-time. To see this, it suffices by Proposition 4.2.4 to find a positive circuit or cocircuit containing a given element e in an orientation \mathcal{O} , which is represented by a matrix obtained from negating corresponding columns of A (by abuse of notation we continue to write as A). e is in some positive circuit of \mathcal{O} if and only if the integer program $\min(\mathbf{1}^T \mathbf{v} : A\mathbf{v} = 0, \mathbf{v}_e = 1, 0 \leq \mathbf{v}_i \leq 1, \mathbf{v}_i \in \mathbb{Z})$ has a solution, and if a solution exists, the support of any minimizer \mathbf{v} is a positive circuit containing e . Since A is totally unimodular, the integer program is actually a linear program, which can be solved in polynomial-time [98]. The cocircuit case is proved analogously.

It remains to show that it is possible to find, in polynomial-time, a coset representative with small (polynomial-size) coefficients in each element of $\text{Jac}(M) \cong \frac{C_1(M)}{\mathbb{Z}_1(M) \oplus B_1(M)}$. For the practical reason of generating random elements of $\text{Jac}(M)$ (cf. Section 5.8.2), we often start with a vector $\mathbf{y} \in \mathbb{Z}^r$ representing a coset of $\frac{\mathbb{Z}^r}{\text{Col}_{\mathbb{Z}}(AA^T)}$, before “lifting” \mathbf{y} to a vector $\gamma \in C_1(M) \cong \mathbb{Z}^E$ that represents an element of $\text{Jac}(M)$. Thus we describe a two-step algorithm to in fact find a “thin” representative in $C_1(M)$ where all coefficients belong to $\{-1, 0, 1\}$ (the existence of which is guaranteed by Proposition 4.2.10), starting with an input vector $\mathbf{y} \in \mathbb{Z}^r$.

In step 1, we first replace \mathbf{y} by $\mathbf{y}' := \mathbf{y} - (AA^T)\lfloor (AA^T)^{-1}\mathbf{y} \rfloor$, where $\lfloor \cdot \rfloor$ is the coordinate-wise truncation. The new vector represents the same element in $\frac{\mathbb{Z}^r}{\text{Col}_{\mathbb{Z}}(AA^T)}$, and it is equal to $(AA^T)((AA^T)^{-1}\mathbf{y} - \lfloor (AA^T)^{-1}\mathbf{y} \rfloor)$. Each coordinate of $(AA^T)^{-1}\mathbf{y} - \lfloor (AA^T)^{-1}\mathbf{y} \rfloor$ is between 0 and 1, and each coordinate of AA^T is between $-m$ and m , so the absolute value of each coordinate of \mathbf{y}' is at most mr . Then we solve the equation $A\gamma = \mathbf{y}'$ to work back in $C_1(M)$, which can be done by choosing an arbitrary basis B and then solving $A|_B\gamma' = \mathbf{y}'$. Since A is totally unimodular, the absolute value of each coefficient of γ' is at most mr^2 .

In step 2, we find a “thin” representative in $[\gamma]$. The procedure described in Proposition 4.2.10 successively chooses a positive (co)circuit C which contains an element e whose coefficient c_e is maximum in γ (recall that we may assume all coefficients in γ are non-negative by reorientation), then subtracts γ_C from γ . An algorithmic optimization is to subtract $\lfloor \frac{c_e}{2} \rfloor \gamma_C$ from γ at once instead. No new element with the absolute value of its coefficient being larger than $\lceil \frac{|\gamma|_{\infty}}{2} \rceil$ is created in each such step, so after every $O(m)$ rounds the maximum absolute value of coefficients is halved, and in a total of $O(m \log m)$ rounds the maximum absolute value of coefficients is reduced to at most 1. \square

Remark. Each of two steps above yields a polynomial-time algorithm by itself. If we only perform the first step, the element γ produced at the end is already of polynomial-size, so we may compute $[\gamma] \cdot [\mathcal{O}]$ using γ directly. However, in such situation we will need to

solve $O(m^2r^2)$ many linear programs to compute the action, while the algorithm in the second step solves $O(m \log m)$ linear programs to find a better γ , and the group action can now be computed by solving only $O(m)$ many extra linear programs. If we perform the second step alone, starting from a general γ as input, the size reducing procedure still terminates in $O(m \log |\gamma|_\infty)$ rounds, so it runs in polynomial-time with respect to the input size.

4.3 Circuit-cocircuit Reversal Systems of Non-Regular Matroids

In [58, Proposition 2], Gioan noticed that Theorem 4.1.4 is not true in general for non-regular oriented matroids. In this section, we prove that Theorem 4.1.4 actually characterizes regular matroids among oriented matroids.

Theorem 4.3.1 *Let M be a non-regular oriented matroid. Then $|\mathcal{G}(M)| < |\mathcal{B}(M)|$.*

4.3.1 Reduction via Circuit-cocircuit Minimal Orientations

Fix a total ordering of the ground set E of M , together with a *reference orientation* of E .

Definition 4.3.2 *An element of E is internally (resp. externally) active in an orientation \mathcal{O} if it is the minimal element in some signed cocircuit (resp. circuit) compatible with \mathcal{O} . The internal (resp. external) activity $\iota(\mathcal{O})$ (resp. $\epsilon(\mathcal{O})$) is the number of internally (resp. externally) active elements in \mathcal{O} .*

We have the following formula for the Tutte polynomial of M using activities of orientations, which was first proven by Las Vergnas.

Theorem 4.3.3 [79] *Let M be an oriented matroid. Then*

$$T_M(x, y) = \sum_{\mathcal{O}} \frac{1}{2^{\iota(\mathcal{O}) + \epsilon(\mathcal{O})}} x^{\iota(\mathcal{O})} y^{\epsilon(\mathcal{O})}, \quad (4.1)$$

where the sum is taken over all $2^{|E|}$ orientations of M .

As immediate corollaries, we have the following enumerations as special evaluations of the Tutte polynomial, extending the list in Proposition 2.5.5.

Corollary 4.3.4 [78] *Let M be an oriented matroid. Then*

1. $T_M(2, 0)$ equals the number of acyclic orientations of M .
2. $T_M(0, 2)$ equals the number of totally cyclic orientations of M .

Now we introduce an important collection of orientations, which will be generalized in the next chapter.

Definition 4.3.5 *An orientation is circuit-cocircuit minimal (with respect to the chosen total ordering and reference orientation of E) if every active element is oriented according to its reference orientation.*

The following results were due to Gioan and Backman, respectively. But in view of their importance in the proof of Theorem 4.3.1, we include a proof for each of them for completeness.

Theorem 4.3.6 [56] *Let M be an oriented matroid. Then the number of circuit-cocircuit minimal orientations of M equals the number of bases of M .*

PROOF: Substituting $x = 1, y = 1$ in (4.1), we have $T_M(1, 1) = \sum_{\iota, \epsilon} \frac{1}{2^{\iota+\epsilon}} |\mathfrak{D}_{\iota, \epsilon}|$, where $\mathfrak{D}_{\iota, \epsilon}$ is the set of orientations with internal activity ι and external activity ϵ . By Proposition 2.5.5, $T_M(1, 1)$ equals the number of bases of M , so it suffices to show that we can partition each $\mathfrak{D}_{\iota, \epsilon}$ into groups of size $2^{\iota+\epsilon}$ such that there is a unique circuit-cocircuit minimal orientation within each group.

Pick an arbitrary orientation \mathcal{O} from $\mathfrak{D}_{\iota, \epsilon}$ (if it is not empty). Let $e_1 < \dots < e_\iota$ (resp. $e'_1 < \dots < e'_\epsilon$) be the elements that are internally (resp. externally) active in \mathcal{O} . For $k = 1, 2, \dots, \iota$, denote by F_k the union of (the support of) all signed cocircuits compatible

with \mathcal{O} whose minimal elements are at least e_k ; dually, for $k = 1, 2, \dots, \epsilon$, denote by F'_k the union of (the support of) all signed circuits compatible with \mathcal{O} whose minimal elements are at least e'_k . The partition $\mathcal{F} = (F_\ell, F_{\ell-1} \setminus F_\ell, \dots, F_1 \setminus F_2; F'_\epsilon, F'_{\epsilon-1} \setminus F'_\epsilon, \dots, F'_1 \setminus F'_2)$ of E is the *active partition* of \mathcal{O} ; note that each component contains precisely one active element.

It is easy to see that reversing the elements of any F_k (resp. F'_k) produces an orientation with the same active partition (hence the same internal and external activities). By induction, reversing any union of components from \mathcal{F} also produces an orientation in $\mathfrak{D}_{\ell, \epsilon}$, and such a reversal procedure induces an equivalence relation on $\mathfrak{D}_{\ell, \epsilon}$. Since each equivalence class contains $2^{\ell+\epsilon}$ orientations, and exactly one orientation in each class is circuit-cocircuit minimal, this gives our desired partition of $\mathfrak{D}_{\ell, \epsilon}$. \square

As a simple corollary, we have the following variation of Theorem 4.3.3.

Corollary 4.3.7 *Let M be an oriented matroid. Then*

$$T_M(x, y) = \sum_{\mathcal{O}} x^{\iota(\mathcal{O})} y^{\epsilon(\mathcal{O})}, \quad (4.2)$$

where the sum is taken over all circuit-cocircuit minimal orientations of M .

Proposition 4.3.8 [5] *Every circuit-cocircuit reversal class of M contains at least one circuit-cocircuit minimal orientation. Furthermore, a circuit-cocircuit minimal orientation equivalent to a given orientation can be obtained greedily.*

PROOF: Start with an arbitrary orientation, and greedily reverse any compatible signed (co)circuit whose minimal element is not oriented according to its reference orientation. Once the procedure stops, we will have a circuit-cocircuit minimal orientation equivalent to the starting orientation, so it suffices to show that the procedure always terminates. If this is not the case, then since the number of orientations is finite, we must (without loss of generality) return to the starting orientation. Let e be the minimal element that was

reversed (which must occur at least twice) in the process. When e was reversed for the first time, we must have reversed it to agree with its reference orientation, so the second reversal is not valid, a contradiction. \square

Remark. We claim that if we always reverse the circuit whose minimal element is the smallest among all “wrongly oriented” positive circuits of the current orientation, then the sequence of minimal elements of the circuits we reverse is strictly increasing, thus we only have to reverse at most m circuits in the greedy process above (and a similar assertion for cocircuits). Suppose that after reversing a positive circuit C of \mathcal{O} whose minimal element is e , we have to reverse a positive circuit C' of the new orientation \mathcal{O}' whose minimal element is $e' < e$. Consider $X := \underline{C} \cup \underline{C}'$, an element in X is either in the positive circuit C of \mathcal{O} or the positive circuit C' of \mathcal{O}' , so by Lemma 4.1.3, \mathcal{O} and \mathcal{O}' restricted to X are totally cyclic. Therefore e' was contained in some positive circuit C'' of \mathcal{O} whose support is contained in X . Since e' is the minimal element of X , e' is also the minimal element of C'' . Furthermore, C'' is necessarily not C as e is already its minimal element, so e' was oriented opposite to its reference orientation in \mathcal{O} already, and we would have reversed C'' instead of C .

Corollary 4.3.9 *The number of circuit-cocircuit reversal classes is at most the number of bases, with equality if and only if no two circuit-cocircuit minimal orientations are contained in the same class.*

As an immediate application of Corollary 4.3.9, we can give a short proof of Theorem 4.1.4: by Lemma 2.3.14, two circuit-cocircuit reversal equivalent orientations of a regular matroid differ by a disjoint union of signed circuits and cocircuits, but for each pair of opposite signed circuits or cocircuits, the minimal element is oriented as its reference orientation in precisely one of them, so at most one orientation within a class can be circuit-cocircuit minimal.

4.3.2 Proof of the Inequality

We prove Theorem 4.3.1 in this section. The idea of our proof is to use the fact that in a non-binary matroid, the symmetric difference of two cocircuits need not be a disjoint union of cocircuits, so the argument (*a posteriori* the conclusion) at the end of Section 4.3.1 fails. We start with an observation on $U_{2,4}$ which can be verified by a slightly tedious but routine case by case analysis.

Proposition 4.3.10 *Up to reorientation, the unique oriented structure on $U_{2,4}$ is the one induced by the four-point line configuration. Namely, the list of signed cocircuits is $(+ + + 0), (+ + 0 -), (+ 0 - -), (0 + + +)$ and their negations.*

Next we construct a special pair of signed cocircuits with respect to a $U_{2,4}$ -minor. These cocircuits will be used in the main proof below.

Lemma 4.3.11 *Let M be an oriented matroid over E and let B be a subset of E such that $M \setminus B \cong U_{2,4} = \{e_1, e_2, e_3, e_4\}$. Without loss of generality the $U_{2,4}$ -minor has the oriented structure described in Proposition 4.3.10. Then there exists a pair of signed cocircuits $\mathbf{D}_1, \mathbf{D}_2$ of M such that*

1. $\mathbf{D}_1(e)\mathbf{D}_2(e) \geq 0$ for every $e \in E$, i.e., \mathbf{D}_1 and \mathbf{D}_2 are conformal.
2. $\mathbf{D}_1(e_1) = \mathbf{D}_1(e_2) = \mathbf{D}_1(e_3) = +, \mathbf{D}_1(e_4) = 0$, and $\mathbf{D}_2(e_1) = \mathbf{D}_2(e_2) = +, \mathbf{D}_2(e_3) \geq 0, \mathbf{D}_2(e_4) = -$.

PROOF: List the elements of B as b_1, \dots, b_m . We will inductively construct a pair of signed cocircuits D_1^i, D_2^i of $M \setminus \{b_{i+1}, \dots, b_m\}$ for each $i = 0, \dots, m$ with the following properties:

1. $D_1^i(e)D_2^i(e) \geq 0$ for every $e \in E \setminus \{b_{i+1}, \dots, b_m\}$.
2. $D_1^i(e_1) = D_1^i(e_2) = D_1^i(e_3) = +, D_1^i(e_4) = 0$ and $D_2^i(e_1) = D_2^i(e_2) = +, D_2^i(e_3) \geq 0, D_2^i(e_4) = -$.

For ease of notation, by reorienting elements in B if necessary, we may assume that each D_1^i is positive and each D_2^i is non-positive outside its common support with D_1^i .

The base case is $D_1^0 = (+ + + 0)$, $D_2^0 = (+ + 0 -)$ as cocircuits of $M \setminus B$.

Suppose we have constructed $D_1^p, D_2^p \subset M \setminus \{b_{p+1}, \dots, b_m\}$ for some p . Consider the single-element extensions D'_1, D'_2 of D_1^p, D_2^p obtained by adding b_{p+1} . Without loss of generality, $D'_1(b_{p+1}) \geq 0$ by reorientation. If $D'_1(b_{p+1})D'_2(b_{p+1}) \geq 0$, then we can take $D_1^{p+1} = D'_1, D_2^{p+1} = D'_2$.

If $D'_1(b_{p+1}) = +, D'_2(b_{p+1}) = -$, apply the strong cocircuit elimination axiom to obtain a cocircuit signed D of $M \setminus \{b_{p+2}, \dots, b_m\}$ such that

$$D^+ \subset (D_1^{'+} \cup D_2^{'+}) \setminus \{b_{p+1}\} = \underline{D_1'} \setminus \{b_{p+1}\}$$

and

$$e_4 \in D^- \subset (D_1'^- \cup D_2'^-) \setminus \{b_{p+1}\} = D_2'^- \setminus \{b_{p+1}\} \subset \underline{D_1'^c}.$$

We claim that choosing $D_1^{p+1} = D'_1$ and $D_2^{p+1} = D$ works. From the containment $D^- \subset \underline{D_1'^c}$, D'_1, D satisfy (1). Also, D'_1 satisfies (2) as the signs of e_1, e_2, e_3, e_4 do not change after the single-element extension. Furthermore, we have $D(e_4) = -$ by construction, and $D(e_1), D(e_2), D(e_3) \geq 0$ by (1), so it remains to show D contains both e_1, e_2 in its support. If not, then by deleting $\{b_1, \dots, b_{p+1}\}$ from D (resp. $M \setminus \{b_{p+2}, \dots, b_m\}$), we have a signed subset X of $\{e_1, \dots, e_4\}$ with $X(e_4) = -$ and otherwise non-negative, but not positive on both e_1, e_2 . By [23, Theorem 3.7.11], X is a signed covector and thus contains a signed cocircuit, but $U_{2,4}$ does not have such a signed cocircuit. \square

PROOF OF THEOREM 4.3.1: By Proposition 2.3.12, M contains $U_{2,4}$ as a minor, say $M/A \setminus B = \{e_1, e_2, e_3, e_4\}$ is isomorphic to $U_{2,4}$, in which we will assume it has the oriented structure described in Proposition 4.3.10. Apply Lemma 4.3.11 to obtain a pair of signed cocircuits $\mathbf{D}_1, \mathbf{D}_2$ of M/A , thus of M itself. By reorientation, we may assume that \mathbf{D}_1 is

non-negative while \mathbf{D}_2 is non-positive outside $\underline{\mathbf{D}}_1 \cap \underline{\mathbf{D}}_2$. Take the positive orientation for each element of E as the reference orientation.

We start with the orientation \mathcal{O}_1 whose all elements are positive. Note that \mathcal{O}_1 is circuit-cocircuit minimal with respect to any ordering of E . \mathbf{D}_1 is compatible with \mathcal{O}_1 , so we can reverse \mathbf{D}_1 in \mathcal{O}_1 . Afterward, $-\mathbf{D}_2$ is compatible with the new orientation by (1) in Lemma 4.3.11, so we can perform a second cocircuit reversal with $-\mathbf{D}_2$. Denote by \mathcal{O}_2 the resulting orientation and by $S = \underline{\mathbf{D}}_1 \triangle \underline{\mathbf{D}}_2$ the set of elements that are negative in \mathcal{O}_2 . It follows that S is in the cocircuit part of \mathcal{O}_2 , and $e_4 \in S \cap \{e_1, e_2, e_3, e_4\} \subset \{e_3, e_4\}$. In particular, $\mathcal{O}_1 \neq \mathcal{O}_2$.

S does not contain a positive cocircuit of \mathcal{O}_2 (and *a priori* does not contain a positive circuit): suppose \mathcal{O}_2 contains a positive cocircuit D whose support is contained in S , by contracting A and deleting B from M , we obtain a covector $D|_{\{e_1, e_2, e_3, e_4\}} \subset \{e_3, e_4\}$ in $U_{2,4}$, which is impossible. Now fix a total ordering of E such that the elements in S are larger than the rest of the elements. Then \mathcal{O}_2 is circuit-cocircuit minimal because any subset whose minimal element is negative in \mathcal{O}_2 would be contained in S , thus not a positive circuit or cocircuit. The theorem follows from Corollary 4.3.9, since we have two distinct circuit-cocircuit minimal orientations $\mathcal{O}_1, \mathcal{O}_2$ that are equivalent. \square

In [58, Theorem 10], several other analogous enumerative results relating the number of circuit/cocircuit reversal classes and special evaluations of Tutte polynomial were given. By restricting ourselves to the circuit or cocircuit part of orientations, as well as further restriction to totally cyclic or acyclic orientations (and their reversal classes), the above argument for Theorem 4.3.1 can be modified to show the non-regular version of those statements.

Proposition 4.3.12 *Let M be a non-regular oriented matroid. Then we have the following.*

1. *The number of acyclic cocircuit reversal classes of M is strictly less than $T_M(1, 0)$.*

2. *The number of totally cyclic circuit reversal classes of M is strictly less than $T_M(0, 1)$.*

3. *The number of cocircuit reversal classes of M is strictly less than $T_M(1, 2)$.*

4. *The number of circuit reversal classes of M is strictly less than $T_M(2, 1)$.*

CHAPTER 5

GEOMETRIC BIJECTIONS

This chapter is mainly based on joint work with Spencer Backman and Matthew Baker [6].

5.1 Statement of the Main Theorem

In this chapter, we study a class of bijections between $\mathcal{B}(M)$ and $\mathcal{G}(M)$ of a regular matroid M . Despite the simple combinatorial description, the only known proofs of their bijectivity rely on geometric combinatorics in an essential way. In order to define the bijection, we introduce several important notions.

Definition 5.1.1 *Let M be a regular (oriented) matroid. A circuit signature σ is map from $\mathcal{C}(\underline{M})$ to $\mathcal{C}(M)$, that is, for every circuit C of the underlying matroid, σ picks out one of the two possible orientations of C . A cocircuit signature is defined similarly.*

A circuit signature is acyclic if $\sum_{C \in \mathcal{C}(\underline{M})} a_C \sigma(C) = 0$ has no non-zero non-negative solution, where each $\sigma(C)$ is considered as a $\{1, 0, -1\}$ -vector in $C_1(M) \cong \mathbb{Z}^E$. An acyclic cocircuit signature is defined similarly.

The notion of acyclic signatures is geometric in nature.

Lemma 5.1.2 *Let σ and σ^* be signatures of $\mathcal{C}(M)$ and $\mathcal{C}^*(M)$, respectively. Then σ and σ^* are both acyclic if and only if there exists $\mathbf{w} \in \mathbb{R}^E$ such that $\mathbf{w} \cdot \sigma(C) > 0$ for every circuit C of \underline{M} and $\mathbf{w} \cdot \sigma^*(D) > 0$ for every cocircuit D of \underline{M} .*

PROOF: The Gordan's alternative [23, p. 478] applied to the circuit-element incidence matrix of M implies σ is acyclic if and only if there exists a vector \mathbf{w} (which can be chosen from the circuit space of M) such that $\mathbf{w} \cdot \sigma(C) > 0$ for every circuit C of M . The

orthogonality of the circuit space and cocircuit space of M allows us to consider circuits and cocircuits simultaneously. \square

Now we define our family of bijections. The following is the main theorem of this chapter, which can be considered as the main theorem of the dissertation as well.

Theorem 5.1.3 *Let M be a regular matroid, and fix acyclic signatures σ and σ^* of $\mathcal{C}(M)$ and $\mathcal{C}^*(M)$, respectively. Given a basis $B \in \mathcal{B}(M)$, let $\mathcal{O}(B)$ be the orientation of M in which we orient each $e \notin B$ according to its orientation in $\sigma(C(B, e))$ and each $e \in B$ according to its orientation in $\sigma^*(C^*(B, e))$. Then the map $B \mapsto [\mathcal{O}(B)]$ gives a bijection $\beta_{\sigma, \sigma^*} : \mathcal{B}(M) \rightarrow \mathcal{G}(M)$.*

Before moving on to the proof of Theorem 5.1.3, we give a few examples of acyclic signatures and a non-example.

Example 5.1.4 *Fix a total ordering and a reference orientation of $E(M)$, and orient each circuit $C \in \mathcal{C}(M)$ compatibly with the reference orientation of the minimal element in C . This gives an acyclic signature of $\mathcal{C}(M)$.*

Indeed, suppose the signature is not acyclic and take some non-trivial expression $\sum_{C \in \mathcal{C}(M)} a_C \sigma(C) = 0$ with $a_C \geq 0$. Let e be the minimal element appearing in some circuit in the support of this expression. Then the element e must be appear with different orientations in at least two different circuits, and thus one of these circuits is not oriented according to σ , a contradiction.

In view of Lemma 5.1.2, we can take $(1, 1/2, 1/4, \dots, 1/2^{m-1})$ as the vector (up to signs of individual coordinates) that induces such acyclic signature.

Example 5.1.5 *Let M be the graphic matroid of a connected graph G . Fix a vertex q of G and a spanning tree T of G , in which we consider q as the root of T . Pick an arbitrary depth first search ordering of the edges in T (or any ordering such that each edge in T is smaller than its descendants), and extend such ordering arbitrarily by declaring all edges outside T*

to be larger than any edge in T . Then we choose the reference orientation of an edge in T as the one orienting away from q (with respect to T), and choose an arbitrary orientation for the remaining edges. By Example 5.1.4, such data induces an acyclic cocircuit signature of M .

We claim that the signature only depends on the choice of q : a cocircuit C of M corresponds to a cut of G , say C is the set of edges between $q \in U \subset V(G)$ and $V(G) \setminus U$, then $C \cap T$ is non-empty, and the minimal edge in the intersection must be oriented from U to $V(G) \setminus U$, hence C is always oriented away from q . We call such cocircuit signature the q -connected signature of M .

Example 5.1.6 *Theorem 5.1.3 is not true for all signatures in general (although the acyclic condition can be weakened, cf. Theorem 7.1.5). Consider the graph in Figure 5.1. Take the cycle signature σ specified by the directed cycles in the top row, and take the v -connected signature σ_v^* as the cut signature. Then the two spanning trees in the bottom row produce the same orientation (class) under the map $\beta_{\sigma, \sigma_v^*}$. Indeed, it is easy to check that the sum of all directed cycles in the top row is zero, hence σ is not acyclic.*

5.2 Zonotopes

We introduce the key geometric notion in the proof of Theorem 5.1.3, we will work in full generality as much as possible unless the total unimodularity condition is necessary.

Definition 5.2.1 *Let A be a (full row rank) $r \times m$ matrix over \mathbb{R} . The column zonotope $Z_A \subset \mathbb{R}^r$ is the Minkowski sum of the columns of A (here we identify a column vector A_i as the line segment from $\mathbf{0}$ to A_i). Let $V^* \leq \mathbb{R}^m$ be the row space (cocircuit space) of A and denote by $\pi_{V^*} : \mathbb{R}^m \rightarrow V^*$ the orthogonal projection.*

The row zonotope \widetilde{Z}_A is the Minkowski sum of the vectors $\pi_{V^}(\mathbf{e}_i)$, $i = 1, 2, \dots, m$, where $\{\mathbf{e}_1, \dots, \mathbf{e}_m\}$ is the standard basis of \mathbb{R}^m .*

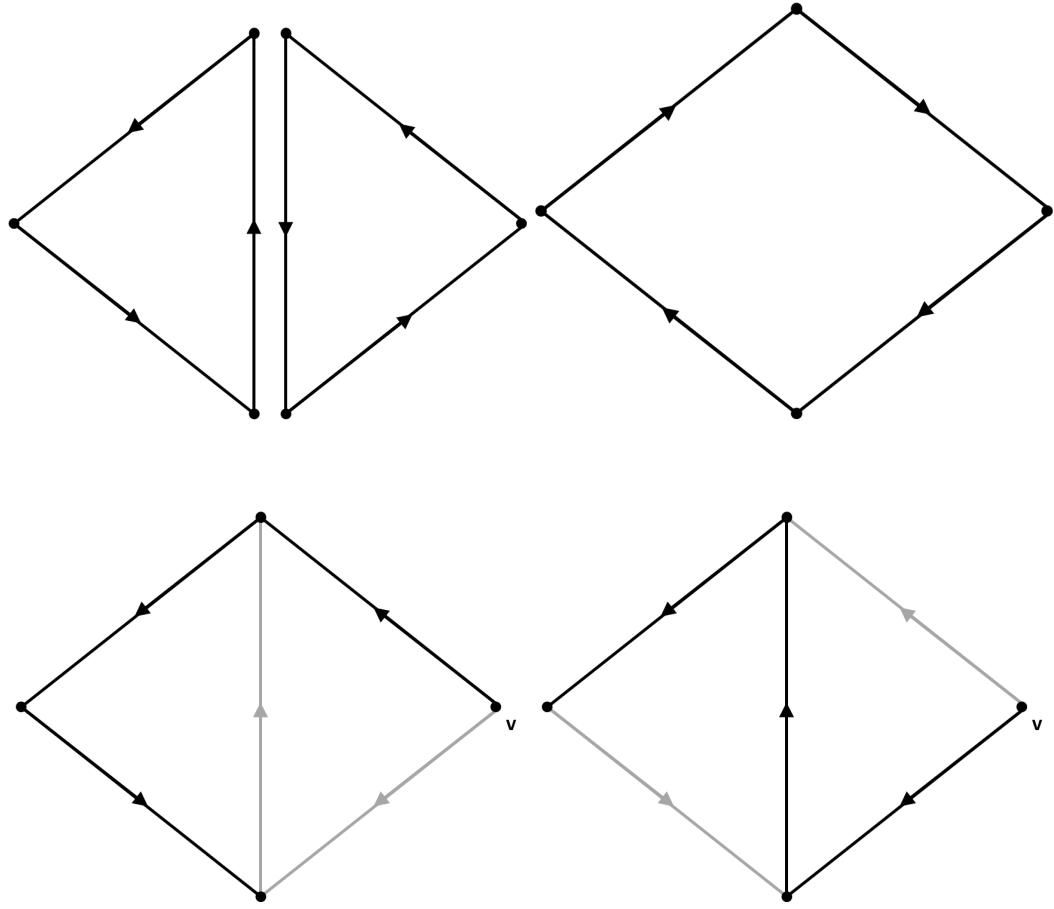


Figure 5.1: Diagrams for Example 5.1.6. The top row specifies a (non-acyclic) cycle signature. The bottom row shows β_{σ, σ^*} need not be bijective when the signatures failed to be acyclic.

It is easy to check that the linear transformation $L : v \mapsto Av$ is a bijection between \widetilde{Z}_A and Z_A , and L preserves the set of lattice points when A is totally unimodular. The reason we introduce two versions of essentially the same object is mostly for the sake of convenience, as one space is often more convenient to work with than the other under different scenarios.

Remark. In the following sections, whenever C is a signed circuit (resp. signed cocircuit) of M , v_C will denote an element of $\ker(A)$ (resp. $\text{row}(A)$) whose sign pattern is C . Such v_C is unique up to multiplication of positive scalars, and in our discussion, either a particular choice will be specified, or an arbitrary choice works. If we have already fixed

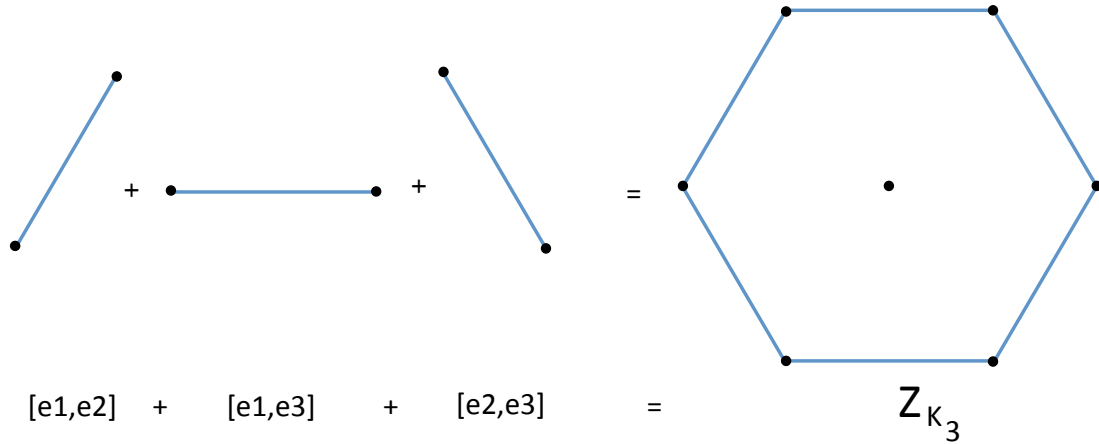


Figure 5.2: The zonotope associated to M_{K_3} . We note that it is living in \mathbb{R}^3 as M_{K_3} is not of full row rank, and we use the line segment from e_i to e_j instead of from 0 to $e_j - e_i$, but none of these differences in convention changes the geometry of zonotope.

some v_C , we automatically choose v_{-C} to be $-v_C$.

Now we relate the zonotope Z_A with the oriented matroid $M := M(A)$ realized by A .

Definition 5.2.2 A continuous orientation \mathcal{O} of M is a function $E \rightarrow [-1, 1]$. If $\mathcal{O}(e) \in \{-1, 1\}$ for all $e \in E$, we say that \mathcal{O} is a discrete orientation.

Let C be a signed circuit of M . A continuous orientation \mathcal{O} is compatible with C if $\mathcal{O}(e) \neq -C(e)$ for all e in the support of C .

Given a continuous orientation \mathcal{O} compatible with a signed circuit C , a continuous circuit reversal with respect to C replaces \mathcal{O} by a new continuous orientation $\mathcal{O} - \epsilon v_C$ for some $\epsilon > 0$; often we choose the maximum of such ϵ , so that at least one $e \in \underline{C}$ satisfies $(\mathcal{O} - \epsilon v_C)(e) = -\text{sign}(C(e))$. The continuous circuit reversal system is the equivalence relation on the set of all continuous orientations of M generated by all possible continuous circuit reversals.

We can make the same definitions for cocircuits analogously.

Immediately from the definition, a (row or column) zonotope is the projection of the

hypercube $[0, 1]^E$ along A , and the hypercube $[0, 1]^E$ can be viewed as the space of all continuous orientations using the relation $\mathcal{O}_{\mathbf{v}}(e) = 2\mathbf{v}_e - 1$ for $\mathbf{v} \in [0, 1]^E$. Therefore we have the following connection between points in zonotope and continuous circuit reversal system.

Proposition 5.2.3 *The map $\psi : [\mathcal{O}] \mapsto \sum_{e \in E} \frac{\mathcal{O}(e)+1}{2} A_e$ gives a bijection between continuous circuit reversal classes of M and points of the zonotope Z_A .*

PROOF: By definition, ψ sends every continuous orientation to some point in Z_A , and ψ is surjective. Since \mathbf{v}_C 's are elements of $\ker(A)$, two continuous orientations in the same circuit reversal class map to the same point of Z_A , so it remains to show the converse. Suppose $\psi(\mathcal{O}) = \psi(\mathcal{O}')$. By Lemma 2.3.9, $\mathcal{O} - \mathcal{O}'$ can be written as a sum of signed circuits in which each signed circuit is compatible with \mathcal{O} , and \mathcal{O} can be transformed to \mathcal{O}' via the corresponding continuous circuit reversals in any order. \square

5.3 Zonotopal Tilings

We explain how an acyclic circuit signature induces a tiling of Z_A by sub-zonotopes (which will be parallelotopes). There are classical approaches to zonotopal tilings that use either polyhedral geometry or abstract oriented matroid theory [23, Section 2.2], [115, Lecture 7]. The approach here was suggested by Backman, and it is more similar to the constructions in tropical divisor theory (e.g. the ABKS decomposition).

Since we are still working at the generality of realizable oriented matroids, we introduce a more general definition of acyclic circuit signatures.

Definition 5.3.1 *Let M be a realizable oriented matroid and let A be a matrix realizing M . A circuit signature is acyclic (with respect to A) if $\sum_{C \in \mathcal{C}(M)} a_C \mathbf{v}_{\sigma(C)} = 0$ has no non-zero non-negative solution. An acyclic cocircuit signature with respect to A is defined similarly.*

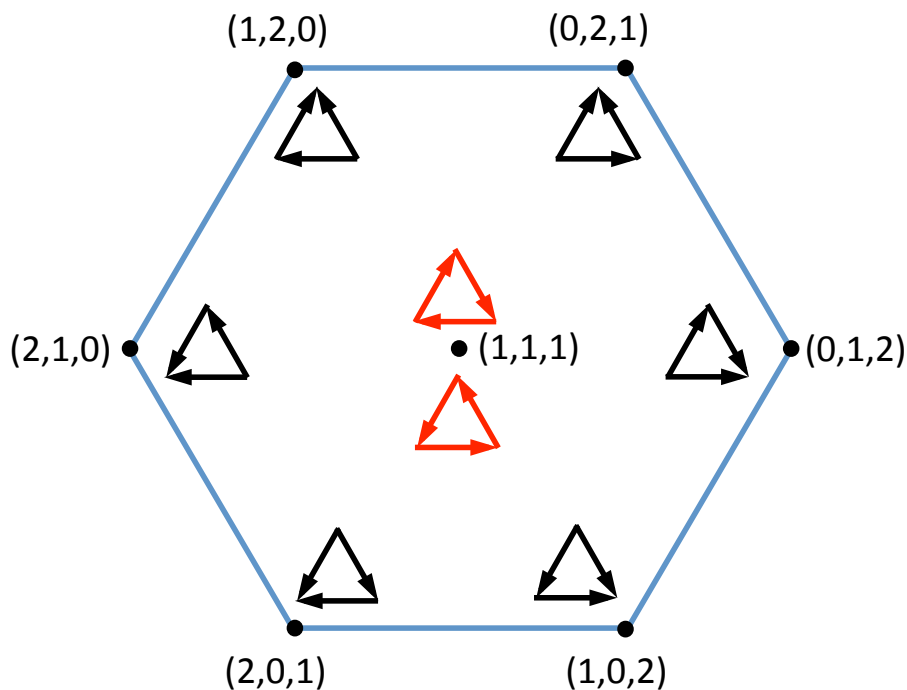


Figure 5.3: The zonotope associated to K_3 , with points identified as continuous circuit reversal classes. The classes corresponding to the vertices of the zonotope are acyclic, hence they are singletons, whereas the lattice point in the middle contains two (discrete) orientations as well as a family of continuous orientations between them.

Remark. When M is a regular matroid represented by a totally unimodular matrix A , $\mathbf{v}_{\sigma(C)}$'s in Definition 5.3.1 can be chosen to be $\{1, 0, -1\}$ -vectors. Furthermore, Proposition 2.3.13 implies the definition is independent of the choice of A . Therefore Definition 5.3.1 is indeed a generalization of Definition 5.1.1.

We have the following straightforward variation of Lemma 5.1.2.

Lemma 5.3.2 *Let σ and σ^* be signatures of $\mathcal{C}(M)$ and $\mathcal{C}^*(M)$, respectively. Then σ and σ^* are both acyclic with respect to A if and only if there exists $\mathbf{w} \in \mathbb{R}^E$ such that $\mathbf{w} \cdot \mathbf{v}_{\sigma(C)} > 0$ for every circuit C of \underline{M} and $\mathbf{w} \cdot \mathbf{v}_{\sigma^*(D)} > 0$ for every cocircuit D of \underline{M} .*

Definition 5.3.3 *Let M be an oriented matroid realized by matrix A . Let σ be an acyclic circuit signature with respect to A . We say a continuous orientation \mathcal{O} is σ -compatible if*

every signed circuit C of M compatible with \mathcal{O} is oriented according to σ .

Theorem 5.3.4 *Let M, A, σ be the same as in Definition 5.3.3. Then each continuous circuit reversal class M contains a unique σ -compatible continuous orientation.*

PROOF: By Lemma 5.3.2, there exists $\mathbf{w} \in \mathbb{R}^E$ such that $\mathbf{w} \cdot \mathbf{v}_{\sigma(C)} > 0$ for every circuit C of M . Consider the function $P(\mathcal{O}') := \mathbf{w} \cdot \mathcal{O}'$. If $-\sigma(C)$ is compatible with \mathcal{O} for some circuit C , then performing a continuous circuit reversal using $-\sigma(C)$ strictly increases the value of P , so every maximizer of P inside a class (if exists) must be σ -compatible. The set of continuous orientations in a continuous circuit reversal class can be identified with the fiber of ψ over a point in Z_A , which is compact as it is a closed subset of the hypercube. Since P is continuous, a maximizer of P , and hence a σ -compatible continuous orientation, must exist in every continuous circuit reversal class.

Now suppose there are two distinct σ -compatible continuous orientations $\mathcal{O}, \mathcal{O}'$ in a continuous circuit reversal class. By Lemma 2.3.9, \mathcal{O} can be transformed to \mathcal{O}' via a series of continuous circuit reversals in which each signed circuit involved is compatible with \mathcal{O} , hence agrees with σ . If the last signed circuit involved in the series of reversals is C , then $-C$ is a signed circuit compatible with \mathcal{O}' . Therefore $-C$ agrees with σ as well, which is a contradiction. This proves the uniqueness of the σ -compatible orientation in each class. \square

Remark. By interpreting σ -compatible orientations as maximizers of the linear function P , it is easy to see that the map $\mu : Z_A \rightarrow [0, 1]^E$, which takes a point z of Z_A to the unique σ -compatible continuous orientation in the continuous circuit reversal class corresponding to z , is a continuous section to the map ψ . Such point of view is closely related to the classical theory of zonotopal tilings.

Next we explain how σ -compatible orientations are relate to bases of M .

Proposition 5.3.5 *Let M, A, σ be the same as in Definition 5.3.3.*

1. *If \mathcal{O} is a σ -compatible continuous orientation, then the set of $e \in E$ which are bi-oriented (i.e., $\mathcal{O}(e) \neq \pm 1$) is independent in M .*

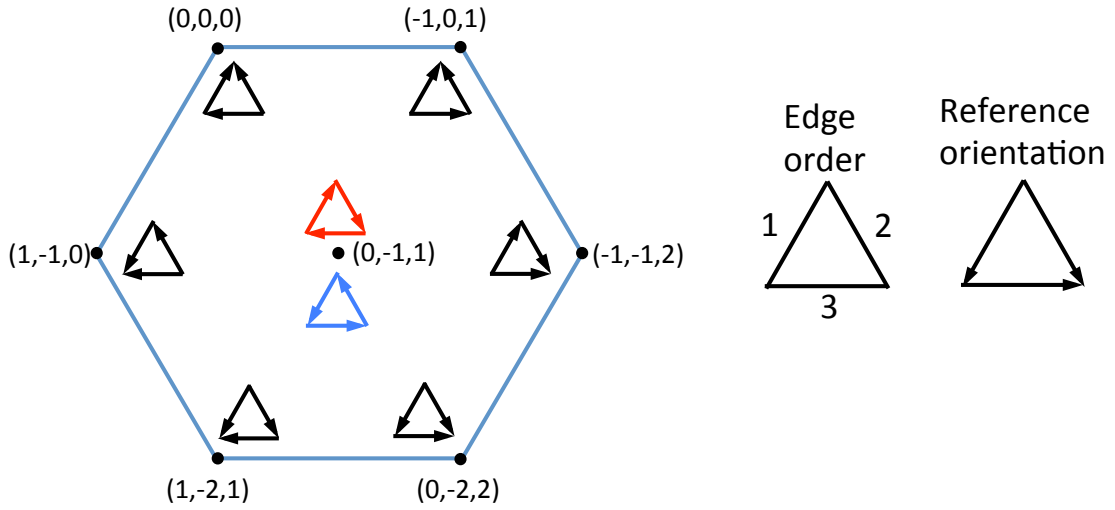


Figure 5.4: The zonotope associated to K_3 with points identified as continuous orientation classes as in Figure 5.3. However, for the lattice point in the middle, we use the σ induced by the total ordering and reference orientation on the right to pick out a σ -compatible (continuous) orientation, which is highlighted in blue.

2. If B is a basis for M and $b : B \rightarrow (-1, 1)$ is any function, then there is a unique σ -compatible continuous orientation $\mathcal{O} = \mathcal{O}(B, b)$ such that $\mathcal{O}(e) = b(e)$ for all $e \in B$ and $\mathcal{O}(e) \in \{\pm 1\}$ for all $e \notin B$.

PROOF: For (1), suppose the set S of bi-oriented elements in a continuous orientation \mathcal{O} is not independent. Then S contains some circuit C , and \mathcal{O} is compatible with both orientations of C , so \mathcal{O} is not σ -compatible.

For the uniqueness assertion in (2), note that each element not in B must be oriented in agreement with the orientation of its fundamental circuit given by σ , as for otherwise the fundamental circuit will be compatible with $-\sigma$. Such unique choice of orientations outside B , together with b itself, gives a continuous orientation \mathcal{O} .

Now we claim that such \mathcal{O} is σ -compatible. If not, then \mathcal{O} is compatible with $-\sigma(C)$ for some circuit C . Without loss of generality, we may assume that $|C \setminus B| \neq 0$ is minimum among all such circuits. Pick any $e \in C \setminus B$ and let C' be the fundamental circuit of

e with respect to B . Then \mathcal{O} is compatible with $\sigma(C')$ by construction. Pick suitable $\mathbf{v}_{-\sigma(C)}, \mathbf{v}_{\sigma(C')}$ such that they agree on the e -th coordinate. Using Lemma 2.3.9, we write $\mathbf{v}_{-\sigma(C)} - \mathbf{v}_{\sigma(C')} = \sum_D \mathbf{v}_D$ with D 's being signed circuits conformal with the left hand side; by Lemma 2.3.9, any D in the sum does not contain e . Since $\mathbf{v}_{\sigma(C)} + \mathbf{v}_{\sigma(C')} + \sum_D \mathbf{v}_D = 0$, at least one such D is oriented opposite to σ by acyclicity. Then we note that such D is compatible with \mathcal{O} : each element of D is either in B (which is bi-oriented in \mathcal{O}), or from C and oriented as in $-\sigma(C)$ (which is compatible with \mathcal{O}). However, $D \setminus B \subset (C \setminus B) \setminus \{e\}$, contradicting the minimality of $|C \setminus B|$. \square

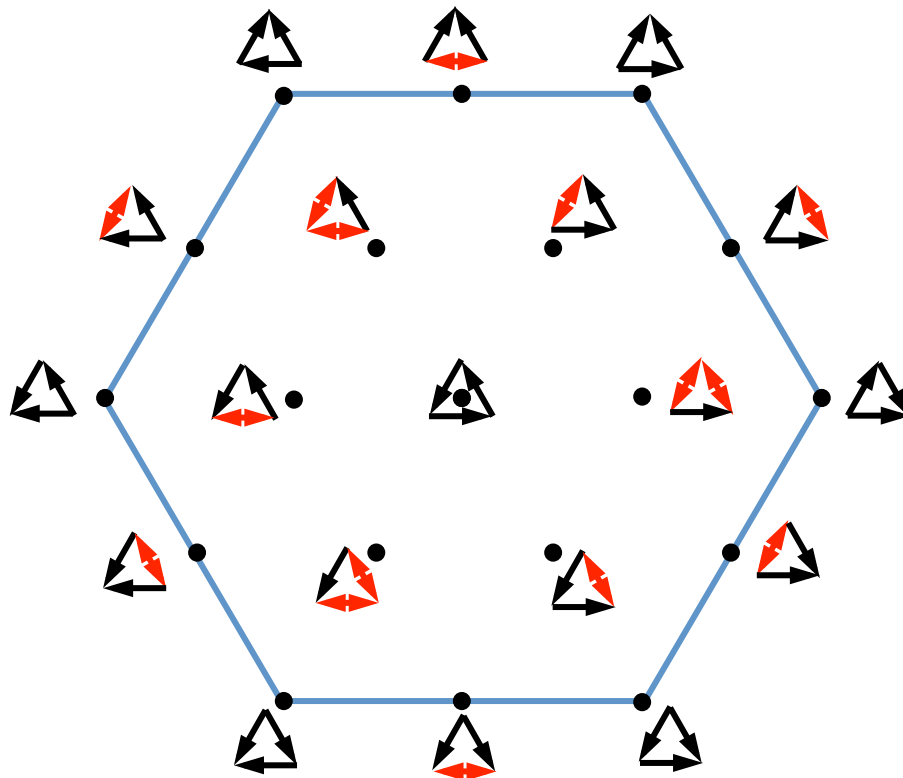


Figure 5.5: More examples of σ -compatible continuous orientations of the zonotope of K_3 . The bi-oriented elements are highlighted in red.

By Proposition 5.3.5, for each basis B of M , every σ -compatible continuous orientation whose bi-oriented part is B has the form $\mathcal{O}(B, b)$ and vice versa. Let $Z^\circ(B)$ be the projection of these continuous orientations to Z_A , and let $Z(B)$ be the topological closure

of $Z^\circ(B)$ in Z_A .

Theorem 5.3.6 1. *The union of $Z(B)$ over all bases B of M is equal to Z_A , and if B, B' are distinct bases then $Z^\circ(B)$ and $Z^\circ(B')$ are disjoint.*

2. *The collection of $Z(B)$ as B varies over all bases B for M gives a polyhedral subdivision of Z_A whose vertices (i.e., 0-cells) correspond via ψ to the σ -compatible discrete orientations of M .*

PROOF: The only non-trivial part is the first half of (1). By Proposition 5.2.3, every point of Z_A is of the form $\psi(\mathcal{O})$ for some continuous orientation \mathcal{O} , and by Theorem 5.3.4 we may assume \mathcal{O} is σ -compatible. Hence by Proposition 5.3.5, it suffices to show that if the set \hat{B} of bi-oriented elements in \mathcal{O} do not form a basis, then we can bi-orient one or more elements in \mathcal{O} while maintaining σ -compatibility. By induction, we will end up with the bi-oriented elements forming a basis B of M , which implies that $\psi(\mathcal{O})$ is a limit point of $Z^\circ(B)$.

Suppose that for every $e \notin \hat{B}$ such that $\hat{B} \cup \{e\}$ is independent in M , bi-orienting e in \mathcal{O} (in an arbitrary way) will cause the new continuous orientation \mathcal{O}_e to no longer be σ -compatible. Then for every such e , \mathcal{O}_e is compatible with $-\sigma(C_e)$ for some circuit C_e of M containing e . Pick, among all such elements e and circuits C_e , the pair that maximizes $w \cdot \mathbf{v}_{\sigma(C_e)}$, where we always choose the “normalized” $\mathbf{v}_{\sigma(C_e)}$ whose e -th coordinate is $\sigma(C_e)(e)$ (which is either 1 or -1).

The circuit C_e must contain another element $f \notin \hat{B}$ such that $\hat{B} \cup \{f\}$ is independent in M , for otherwise e is contained in the closure of \hat{B} . By assumption there exists some circuit C_f containing f such that \mathcal{O}_f is compatible with $-\sigma(C_f)$. The signs of $\sigma(C_e)$ and $\sigma(C_f)$ over f are different, so we can choose a suitable positive multiple \mathbf{v}_f of $\mathbf{v}_{\sigma(C_f)}$ such that the f -th coordinates of \mathbf{v}_f and $\mathbf{v}_{-\sigma(C_e)}$ are equal.

By a trivial modification of Lemma 2.3.9, $\mathbf{v}_{-\sigma(C_e)} - \mathbf{v}_f$ can be written as a weighted sum $\sum_{i=1}^k \lambda_i \mathbf{v}_{C_i}$ of signed circuits C_i 's with positive λ_i 's (we introduce λ_i 's because we

will specify the vectors \mathbf{v}_{C_i} 's explicitly, and in general our choices only coincide with the vectors provided by Lemma 2.3.9 up to multiplication of positive scalars). Each such signed circuit C_i that does not contain e must be compatible with \mathcal{O} (hence σ), while those signed circuits that contain e would at least be compatible with \mathcal{O}_e . Since $w \cdot (\sum_{i=1}^k \lambda_i \mathbf{v}_{C_i}) = w \cdot (\mathbf{v}_{-\sigma(C_e)} - \mathbf{v}_f) < 0$, some C_i 's are not compatible with σ , which implies that they are not compatible with \mathcal{O} , thus they contain e . In particular, the assertion of Lemma 2.3.9 guarantees that the sign of the e -th coordinate of $\mathbf{v}_{-\sigma(C_e)} - \mathbf{v}_f$ agrees with $-\sigma(C_e)$. But as the signs of $\sigma(C_e)(e)$ and $\sigma(C_f)(e)$ are different, the absolute value of the e -th coordinate is at most the absolute value of the e -th coordinate of $\mathbf{v}_{-\sigma(C_e)}$, which is 1.

Without loss of generality, the circuits containing e are C_1, C_2, \dots, C_j . We choose \mathbf{v}_{C_i} 's so that their e -th coordinates equal $-\sigma(C_e)(e)$. By comparing e -th coordinate, $\sum_{i=1}^j \lambda_i \leq 1$. Now we have

$$w \cdot \left(\sum_{i=1}^j \lambda_i \mathbf{v}_{C_i} \right) = w \cdot \left(\mathbf{v}_{-\sigma(C_e)} - \mathbf{v}_f - \sum_{i=j+1}^k \lambda_i \mathbf{v}_{C_i} \right) < -w \cdot \mathbf{v}_{\sigma(C_e)} < 0,$$

i.e., there exists some C_i with $1 \leq i \leq j$ that is compatible with \mathcal{O}_e , disagrees with σ , and $w \cdot \mathbf{v}_{\sigma(C_i)} > w \cdot \mathbf{v}_{\sigma(C_e)}$, contradicting our choice of C_e . \square

From Theorem 5.3.6, $Z(B)$'s subdivide Z_A . Moreover, each $Z(B)$ can be identified with the parallelotope $[0, 1]^B$, where the e -th coordinate of a point in $Z(B)$ (corresponding to a σ -compatible continuous orientation \mathcal{O} via ψ) is the value $\frac{\mathcal{O}(e)+1}{2}$. Therefore restricting to a face of $Z(B)$ of codimension i can be described as orienting the corresponding i elements in B .

Finally, we describe the incidence relations of the cells $Z(B)$'s. Here we fix an acyclic circuit signature σ that induces the tiling.

Proposition 5.3.7 *Let B be a basis, let $e \in B$, and let \mathcal{O} be any continuous orientation of the form $\mathcal{O}(B, b)$ (cf. Proposition 5.3.5). Choose an orientation for e , and let \mathcal{O}' be the continuous orientation obtained from \mathcal{O} by orienting e according to our choice. From the*

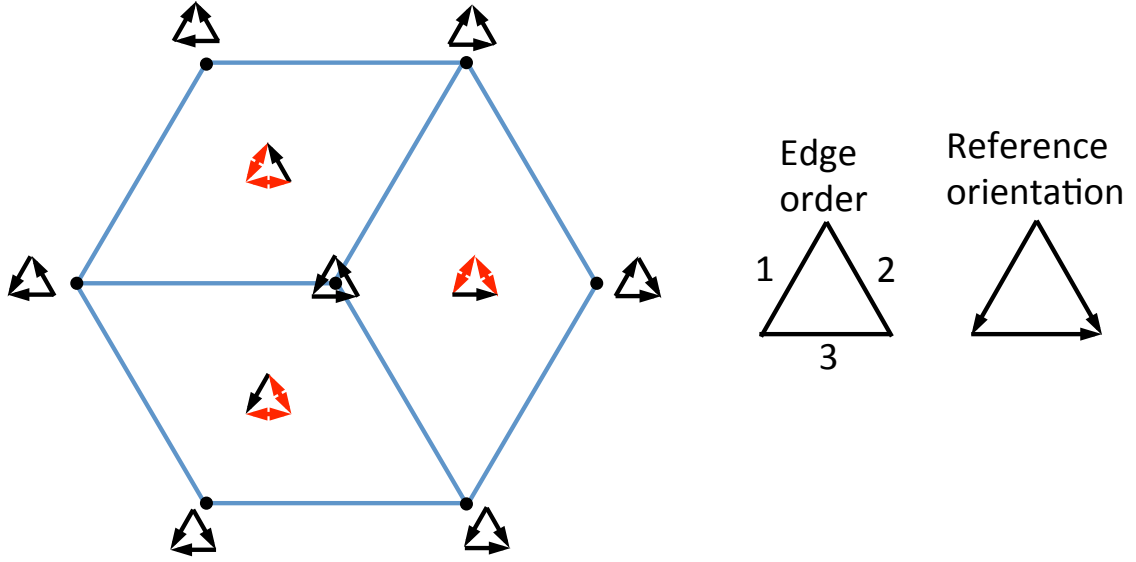


Figure 5.6: The subdivision of the zonotope associated to K_3 as described in Theorem 5.3.6 using σ induced by the total ordering and reference orientation on the right. The red edges are bi-oriented.

remark above, \mathcal{O} is in the interior of $Z(B)$ and \mathcal{O}' is in the relative interior of a facet F_e of $Z(B)$. Denote by K the (signed) fundamental cocircuit of e with respect to B , where we choose the signed version that agrees with e 's chosen orientation. Then either

1. \mathcal{O}' is compatible with K , in which case F_e is a facet of Z_A , or
2. there exists a unique element $f \in K \setminus \{e\}$ such that the orientation obtained by reversing f in \mathcal{O}' is also σ -compatible, and F_e is incident to $Z(B \setminus \{e\} \cup \{f\})$.

PROOF: For simplicity, we use the linear transformation L introduced after Definition 5.2.1 and work with \widetilde{Z}_A . Suppose \mathcal{O}' is compatible with K . Then \mathcal{O}' , as point in \widetilde{Z}_A , is a maximizer of the linear functional $\mathbf{w} \mapsto \mathbf{v}_K \cdot \mathbf{w}$ over \widetilde{Z}_A , so F_e is a facet of \widetilde{Z}_A .

If \mathcal{O}' is not compatible with K , then there exists some element $f \in K$ whose orientation in \mathcal{O}' is inconsistent with e with respect to K . Let C_f be the fundamental circuit of f with respect to B . Since $\pi_{V^*}(\mathbf{v}_{C_f}) = 0$, to move outside of $\widetilde{Z}(B)$ from $\mathcal{O}' \in \widetilde{F}_e$, along the direction which we have used to move from \mathcal{O} to \mathcal{O}' , is the same as bi-orienting the other

elements of C_f in \mathcal{O}' . All elements in C_f except e and f are still bi-oriented in \mathcal{O}' , so the only combinatorial change here is to bi-orient f as long as the moving distance is sufficiently small. As a result, in order to prove (2), we only need to show that for one and only one such element f , bi-orienting f in \mathcal{O}' makes the new orientation \mathcal{O}_f σ -compatible.

Uniqueness: Suppose the continuous orientations $\mathcal{O}_f, \mathcal{O}_g$, obtained by bi-orienting f and bi-orienting g in \mathcal{O}' , respectively, are both σ -compatible. Consider the fundamental circuit C of g with respect to $B \setminus \{e\} \cup \{f\}$, which is also the fundamental circuit of f with respect to $B \setminus \{e\} \cup \{g\}$. Since both f and g are oriented opposite to K in \mathcal{O}' and $K \cap C = \{f, g\}$, \mathcal{O}_f and \mathcal{O}_g would be compatible with opposite orientations of C by orthogonality of C and K , contradicting with the assumption that both \mathcal{O}_f and \mathcal{O}_g are σ -compatible.

Existence: First we prove that if \mathcal{O}_f is not σ -compatible, then the fundamental circuit C_g^f of some $g \in K$ with respect to $B \setminus \{e\} \cup \{f\}$ is not compatible with σ . Intuitively, this is to say that if \mathcal{O}_f is not σ -compatible, we can find a “certificate” by just checking the fundamental circuits of elements in K with respect to $B \setminus \{e\} \cup \{f\}$.

Suppose \mathcal{O}_f is compatible with $-\sigma(C)$ for some circuit C . For each fundamental circuit C_g^f of the element g with respect to $B \setminus \{e\} \cup \{f\}$, we pick the orientation that is compatible with \mathcal{O}_f . Then we have $\mathbf{v}_{-\sigma(C)} = \sum_{g \in C \setminus (B \setminus \{e\} \cup \{f\})} \mathbf{v}_{C_g^f}$ for some suitable choice of $\mathbf{v}_{C_g^f}$'s. Since σ is acyclic, at least one of the C_g^f 's is not compatible with σ . Such g must be from K , or otherwise \mathcal{O}' would have been compatible with such C_g^f already.

Now we apply a greedy procedure to find (the unique) f . Pick any $f_1 \in K$ whose orientation in \mathcal{O}' is inconsistent with e (with respect to K). If \mathcal{O}_{f_1} is σ -compatible then we are done, otherwise there must exist some other $f_2 \in K$ such that $C_{f_2}^{f_1}$ is not compatible with σ . Consider \mathcal{O}_{f_2} next, proceed until the correct f is found. We claim that the procedure always terminates. Suppose not, without loss of generality we may assume $f_{r+1} = f_1$ for some r . By inductively choosing suitable scalar multiples to cancel out the f_i -coordinate

in the partial sum, the support of $\mathbf{v}_{C_{f_2}^{f_1}} + \mathbf{v}_{C_{f_3}^{f_2}} + \dots + \mathbf{v}_{C_{f_1}^{f_r}}$ is contained in a basis, which means the sum is 0 as it is also an element from the circuit space. Such equality contradicts the acyclicity assumption on σ . \square

5.4 Proof of the Main Theorem

In this section, we prove Theorem 5.1.3 using the geometric tools we have developed. We will prove a variation of Theorem 5.1.3 (Theorem 5.4.2) for all realizable oriented matroids, and then specialize it to Theorem 5.1.3 when the oriented matroid is regular.

5.4.1 Proof for Realizable Oriented Matroids

Definition 5.4.1 *Let M be an oriented matroid realized by matrix A . Let σ (resp. σ^*) be an acyclic circuit (resp. cocircuit) signature with respect to A . We say a continuous orientation \mathcal{O} is (σ, σ^*) -compatible if every signed circuit (resp. cocircuit) of M compatible with \mathcal{O} is oriented according to σ (resp. σ^*). Denote by $\chi(M; \sigma, \sigma^*)$ be the set of (σ, σ^*) -compatible discrete orientations.*

Theorem 5.4.2 *Let M, A, σ, σ^* be the same as in Definition 5.4.1. Given a basis $B \in \mathcal{B}(M)$, let $\mathcal{O}(B)$ be the orientation of M in which we orient each $e \notin B$ according to its orientation in $\sigma(C(B, e))$ and each $e \in B$ according to its orientation in $\sigma^*(C^*(B, e))$. Then the map $B \mapsto \mathcal{O}(B)$ gives a bijection $\hat{\beta}_{\sigma, \sigma^*} : \mathcal{B}(M) \rightarrow \chi(M; \sigma, \sigma^*)$.*

The strategy to prove Theorem 5.4.2 (hence Theorem 5.1.3) is to show that the map $\hat{\beta}$ (resp. β) can be thought as a “shifting” map that matches vertices and cells in the zonotopal tiling induced by σ . Recall that by Lemma 5.3.2, there exists a vector \mathbf{w} inducing σ^* , which can be assumed to be in the row space of A by projecting if necessary. Since we are identifying the index set $\{1, 2, \dots, m\}$ of \mathbb{R}^m with elements of E , for an element $e \in E$, we write \mathbf{e} the unit vector whose e -th coordinate is 1.

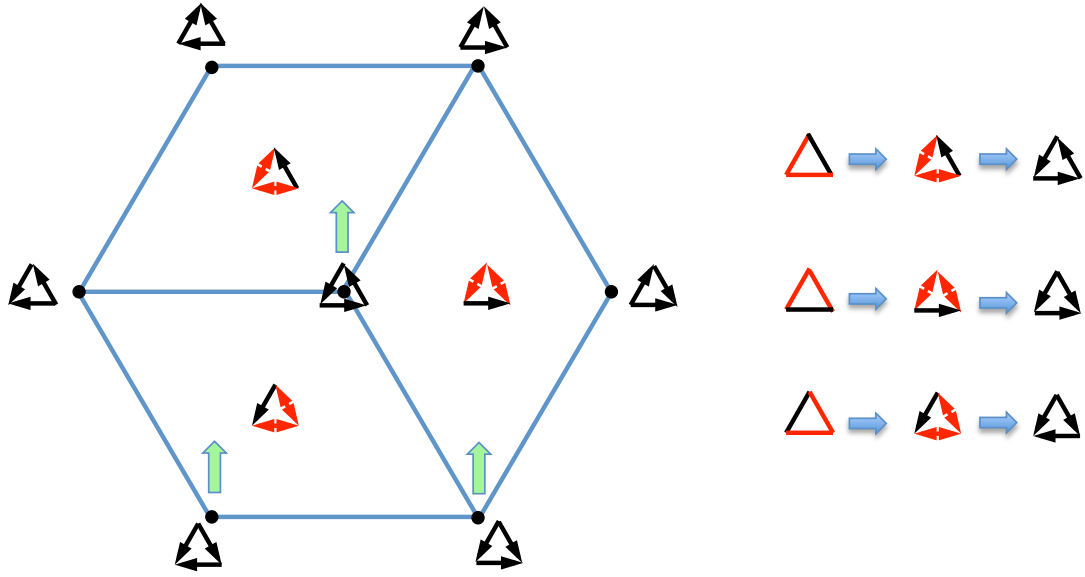


Figure 5.7: Geometric picture of the bijection for K_3 using the pair (σ, σ^*) induced by the total ordering and reference orientation from Figure 5.6.

Lemma 5.4.3 *Let B be a basis of M and let $\widetilde{Z}(B)$ be the cell corresponding to B in the zonotopal tiling of \widetilde{Z}_A induced by σ . Then for all sufficiently small $\epsilon > 0$, the image of $\widetilde{Z}(B)$ under the map $\mathbf{v} \mapsto \mathbf{v} + \epsilon \mathbf{w}$ contains a unique point corresponding (via ψ) to a σ -compatible (discrete) orientation \mathcal{O}_B .*

PROOF: By Theorem 5.3.6, the vertices of $\widetilde{Z}(B)$ correspond to σ -compatible (discrete) orientations. It therefore suffices to prove that \mathbf{w} does not lie in the affine span of any facet of $\widetilde{Z}(B)$.

The affine span of a facet of $\widetilde{Z}(B)$ is spanned by directions of the form $\pi_{V^*}(\mathbf{e})$ for $e \in \hat{B}$, where $\hat{B} \subsetneq B$ is a proper subset of B of size $r - 1$, and we know there is a cocircuit K of M avoiding \hat{B} . Any direction $\mathbf{w}' := \sum_{e \in \hat{B}} \lambda_e \pi_{V^*}(\mathbf{e})$ in the span satisfies $\mathbf{w}' \cdot \mathbf{v}_{\sigma^*(K)} = \sum_{e \in \hat{B}} \lambda_e \mathbf{e} \cdot \mathbf{v}_{\sigma^*(K)} = 0$, since π_{V^*} is self-adjoint and $\pi_{V^*}(\mathbf{v}_{\sigma^*(K)}) = \mathbf{v}_{\sigma^*(K)}$. On the other hand, since \mathbf{w} induces σ^* , we have $\mathbf{w} \cdot \mathbf{v}_{\sigma^*(K)} > 0$. \square

We then define ϕ to be the map that takes a basis B to the orientation \mathcal{O}_B defined in

Lemma 5.4.3.

Theorem 5.4.4 *The map ϕ coincides with the previously defined map $\hat{\beta}$.*

PROOF: Let B be a basis. Then $\phi(B)$ can be obtained by orienting each (bi-oriented) basis element from a continuous σ -compatible orientation in the interior of $\widetilde{Z}(B)$ (which is of the form $\mathcal{O}(B, b)$), so by the greedy procedure described in Proposition 5.3.5, the elements outside B are oriented according to their fundamental circuits, hence $\phi(B)$ agrees with $\hat{\beta}(B)$ outside B .

For elements inside B , we work with the basis $\{\pi_{V^*}(\mathbf{e}) : e \in B\}$ for $V^*(M)$ and write $\mathbf{w} = \sum_{e \in B} w_e \pi_{V^*}(\mathbf{e})$. Identify $\widetilde{Z}(B)$ with $[0, 1]^B$ and the vertices of $\widetilde{Z}(B)$ with $\{0, 1\}^B$. If a vertex v is identified with $(s_e : e \in B)$, $s_e \in \{0, 1\}$, then it corresponds to a σ -compatible discrete orientation where each element $e \in B$ is oriented in agreement with (resp. opposite to) its reference orientation when $s_e = 1$ (resp. $s_e = 0$). The cell $\widetilde{Z}(B)$ will contain v after shifting if and only if the sign pattern of the s_e 's agrees with the sign pattern of the w_e 's, that is, if and only if $s_e = 1$ precisely when $w_e > 0$.

Let $f \in B$, and let K be the fundamental cocircuit of f with respect to B . Then $\sigma^*(K)$ is by definition the orientation of K with the property that $\mathbf{w} \cdot \mathbf{v}_{\sigma^*(K)} > 0$. By a calculation similar to the above,

$$\mathbf{w} \cdot \mathbf{v}_{\sigma^*(K)} = \sum_{e \in B} w_e \mathbf{e} \cdot \mathbf{v}_{\sigma^*(K)} = w_f \mathbf{f} \cdot \mathbf{v}_{\sigma^*(K)},$$

as f is the unique element in $B \cap K$. If $w_f > 0$, then $\mathbf{f} \cdot \mathbf{v}_{\sigma^*(K)} > 0$ and the reference orientation of f agrees with $\sigma^*(K)$, i.e., the orientation of f in $\phi(B)$ is the same as the reference orientation of f . From the last paragraph, f is oriented according to its reference orientation in $\hat{\beta}(B)$ as well, because $w_f > 0$. A similar analysis in the case where $w_f < 0$ yields the same conclusion that $\phi(B)(f) = \hat{\beta}(B)(f)$. \square

Proposition 5.4.5 *Let B be a basis. Then $\hat{\beta}(B)$ is (σ, σ^*) -compatible.*

PROOF: Since $\phi(B)$ is σ -compatible, $\hat{\beta}(B)$ is σ -compatible as well by Theorem 5.4.4. And since the procedure described in Theorem 5.4.2 is symmetric with respect to circuits and cocircuits, a dual argument shows that $\hat{\beta}(B)$ is σ^* -compatible. \square

PROOF OF THEOREM 5.4.2: By Proposition 5.4.5, $\hat{\beta}$ is well-defined. It is injective for the simple geometric reason that a vertex can only be contained in the interior of at most one cell $\widetilde{Z}(B)$ after shifting. To prove surjectivity, we need to show that for every (σ, σ^*) -compatible orientation \mathcal{O} , there exists a continuous orientation \mathcal{O}' such that the displacement from \mathcal{O}' to \mathcal{O} , interpreted as points of \widetilde{Z}_A via ψ , is \mathbf{w} (here we assume \mathbf{w} is sufficiently short). For simplicity, we negate suitable columns of A in order to assume without loss of generality that $\mathcal{O} \equiv \mathbf{1}$, and we modify \mathbf{w} accordingly. For such to be determined \mathcal{O}' , denote by $\mathbf{f}_e \geq 0$ the difference between 1 and $\frac{\mathcal{O}'(e)+1}{2}$. By an easy application of the formula for orthogonal projection, our condition on \mathcal{O}' in terms of displacement becomes $A\mathbf{f} = A\mathbf{w}$, hence \mathcal{O}' exists if and only if the polyhedron

$$\{\mathbf{f} : A\mathbf{f} = A\mathbf{w}, \mathbf{f} \geq \mathbf{0}\} \quad (5.1)$$

is non-empty. But the σ^* -compatible condition implies “if $\mathbf{z}^T A \geq \mathbf{0}$ (i.e., $\mathbf{z}^T A$ is a non-negative sum of signed cocircuits), then $(\mathbf{z}^T A)\mathbf{w} \geq 0$ ”, which is the same as “there exists no \mathbf{z} such that $\mathbf{z}^T A \geq \mathbf{0}$, $\mathbf{z}^T(A\mathbf{w}) < 0$ ”, by the Farkas lemma, the latter condition is equivalent to the existence of some $\mathbf{f} \geq \mathbf{0}$ such that $A\mathbf{f} = A\mathbf{w}$. \square

5.4.2 Specialization to Regular Matroids

In this section, we assume M is a regular (oriented) matroid realized by a totally unimodular matrix A , and we will switch back to the discrete setting from Section 4.1 and 5.1. In particular, we give the discrete version of Definition 5.3.3.

Definition 5.4.6 *Let σ, σ^* be an acyclic circuit signature and an acyclic cocircuit signa-*

ture, respectively. An orientation \mathcal{O} is σ -compatible if every signed circuit C of M compatible with \mathcal{O} is oriented according to σ ; \mathcal{O} is (σ, σ^*) -compatible if furthermore every signed cocircuit compatible with \mathcal{O} is oriented according to σ^* .

Proposition 5.4.7 *The map ψ induces a bijection between circuit reversal classes of M and lattice points of Z_A .*

PROOF: As the columns of A are integral, ψ takes an orientation of M to a lattice point of Z_A ; conversely, for any lattice point $\mathbf{y} \in Z_A \cap \mathbb{Z}^r$ we know $A\hat{\alpha} = \mathbf{y}$, $0 \leq \hat{\alpha}_i \leq 1 \forall i$ has a solution $\hat{\alpha}$, but by the total unimodularity of A , $\hat{\alpha}$ can be chosen to be integral and hence corresponds to an orientation. Thus the image of ψ is precisely the set of lattice points of Z_A . By the orthogonality of circuit space and cocircuit space, two orientations in the same circuit reversal class map to the same point of Z_A . Conversely, suppose $\psi(\mathcal{O}) = \psi(\mathcal{O}')$. By Lemma 2.3.14, $\mathcal{O} - \mathcal{O}'$ can be written as a sum of disjoint signed circuits in which each signed circuit is compatible with \mathcal{O} , and \mathcal{O} can be transformed to \mathcal{O}' via the corresponding circuit reversals in any order. \square

Proposition 5.4.8 *Each circuit reversal class M contains a unique σ -compatible orientation.*

PROOF: If a circuit reversal class contains two distinct σ -compatible orientations $\mathcal{O}, \mathcal{O}'$, then by Lemma 2.3.14, $\mathcal{O}, \mathcal{O}'$ differ by a disjoint union of signed circuits. But for any such circuit C , \mathcal{O} and \mathcal{O}' are compatible with different orientations of C , a contradiction.

For existence, start with any orientation \mathcal{O} in the class and reverse some signed circuit C compatible with \mathcal{O} but not compatible with σ . We claim that the process will eventually stop. Suppose not, since the number of discrete orientations of M is finite, the orientation will without loss of generality return to \mathcal{O} after reversing some signed circuits C_1, \dots, C_k in that order (the circuits might not be distinct). Then $C_1 + \dots + C_k = 0$, which means that $\sigma(C_1) + \dots + \sigma(C_k) = 0$, contradicting the acyclicity of σ . \square

Corollary 5.4.9 *The lattice points of Z_A are exactly the vertices of the zonotopal tiling induced by σ .*

PROOF: By Theorem 5.3.6, the vertices of the zonotopal tiling are precisely the σ -compatible discrete orientations, which form a set of representatives for the (discrete) circuit reversal classes by Proposition 5.4.8, and finally by Proposition 5.4.7, they correspond to the lattice points of Z_A . \square

Proposition 5.4.10 *Every circuit-cocircuit reversal class of M contains a unique (σ, σ^*) -compatible orientation.*

PROOF: By Proposition 4.1.1 and Lemma 4.1.3, it suffices to consider the circuit (resp. cocircuit) part of orientations inside a circuit-cocircuit reversal class, in which the statement follows from Proposition 5.4.8. \square

PROOF OF THEOREM 5.1.3: β_{σ, σ^*} is the composition of $\hat{\beta}_{\sigma, \sigma^*}$ and the map $\mathcal{O} \mapsto [\mathcal{O}]$, which are both bijective by Theorem 5.4.2 and Lemma 5.4.10, respectively. \square

5.5 Ehrhart Theory

We start with the definition of *Ehrhart polynomial* associated to a *lattice polytope*, that is, a polytope whose vertices have integer coordinates.

Theorem 5.5.1 [48] *Let P be a lattice polytope. Then there exists a polynomial E_P (the Ehrhart polynomial) such that, for every positive integer q , $E_P(q)$ equals the number of lattice points in the q -fold dilation of P .*

Now let M be a regular matroid. By Proposition 2.3.13, the zonotopes Z_A 's defined by different totally unimodular realizations A of M differ by lattice preserving full rank linear transformations, so in terms of Ehrhart theory, it is well-defined to talk about the zonotope Z_M associated to M . Stanley's formula relates the Ehrhart polynomial of Z_M and the Tutte polynomial of M .

Theorem 5.5.2 [102] *Let M be a regular matroid. Then $E_{Z_M}(q) = q^r T_M(1 + \frac{1}{q}, 1)$.*

We have the following theorem relating special orientations (resp. orientation classes), special lattice points of Z_M , and special evaluations of the Tutte polynomial.

Theorem 5.5.3 *For a regular matroid M , we have:*

1. *The lattice points of Z_M correspond to the circuit reversal classes of M , and there are $T_M(2, 1)$ many of them.*
2. *The interior lattice points of Z_M correspond to the totally cyclic circuit reversal classes of M , and there are $T_M(0, 1)$ many of them.*
3. *The vertices of Z_M correspond to the acyclic circuit reversal classes of M (hence the acyclic orientations themselves), and there are $T_M(2, 0)$ many of them.*
4. *The volume of Z_M equals the number of bases of M , which is $T_M(1, 1)$.*

PROOF: The first correspondence is Proposition 5.4.7. Next, we know from the proof of Proposition 5.3.7 that an orientation contains a signed cocircuit if and only if it is on the boundary, this shows the second correspondence. For the third correspondence, note that each vertex is contained in r linearly independent facets of Z_M , hence by Proposition 5.3.7, the corresponding orientation contains r independent (in the sense of cocircuit space) signed cocircuits, so it is acyclic (for otherwise we can contract the circuit part of the orientation to obtain a matroid of smaller rank, which still contains r independent signed cocircuits). The enumerative claims concerning orientation classes can be found in [5, 7].

For the last statement, fix an arbitrary acyclic circuit signature σ and consider the zonotopal tiling induced by σ . The number of cells equals the number of bases of M , while each cell has volume 1 by the total unimodularity of A . □

Corollary 5.5.4 *For any acyclic circuit signature σ of a regular matroid M , the number of σ -compatible orientations of M equals $T_M(2, 1)$.*

PROOF: It follows from the first statement of Theorem 5.5.3 and Proposition 5.4.8. \square

We now prove a finer version of Theorem 5.5.2 using the language of orientations. To do so, we first review the notion of *partial orientations* (or (type B) fourorientations in the original paper of Backman and Hopkins [7]) of an oriented matroid.

Definition 5.5.5 *Let M be an oriented matroid. A partial orientation \mathcal{O} assigns to each element of M a value of $+$, $-$, or 0 . Here we say an element is bi-oriented in \mathcal{O} if $\mathcal{O}(e) = 0$.*

For M regular, we can define compatibility of a signed circuit of M with respect to a partial orientation analogous to Definition 5.2.2, and from there we can also define circuit reversals and σ -compatibility of partial orientations (with respect to some acyclic circuit signature σ of M) analogous to Definition 4.1.2 and Definition 5.3.3, respectively.

Lemma 5.5.6 *For any acyclic circuit signature σ of a regular matroid M , there exists a unique σ -compatible partial orientation within each circuit reversal class of partial orientations, in which the bi-oriented part of any σ -compatible partial orientation is an independent set of M .*

PROOF: The proof is analogous to Proposition 5.4.8 and the first half of Proposition 5.3.5. \square

Theorem 5.5.7 *Let σ be an acyclic circuit signature of a regular matroid M and let S be an independent set of M . Then the number of σ -compatible partial orientations whose bi-oriented part is precisely S is $T_{M/S}(2, 1)$.*

We first prove an auxiliary lemma. Recall that given a subset S of a matroid M , its *closure* \bar{S} is the unique maximal subset containing S such that the rank of \bar{S} equals the rank of S . It follows that $e \notin S$ is in \bar{S} if and only if $S \cup \{e\}$ contains a unique circuit, in which we will say it is the fundamental circuit of e (with respect to S) as well.

Lemma 5.5.8 *Let σ be an acyclic circuit signature of a regular matroid M and let S be an independent set of M . Let \mathcal{O} be a partial orientation of M whose bi-oriented part is precisely S , and every element $e \in \overline{S} \setminus S$ is oriented according to the orientation of its fundamental circuit specified by σ . If \mathcal{O} is not σ -compatible, then there exists a signed circuit $C \subset S \cup (M \setminus \overline{S})$ that is compatible with \mathcal{O} but not σ .*

PROOF: The proof is analogous to Proposition 5.3.5. We pick, among all signed circuits that are compatible with \mathcal{O} but not σ , a C such that $|C \setminus S|$ is minimum. If there exists some element $e \in C \cap (\overline{S} \setminus S)$, then we can apply the circuit elimination axiom to C and the fundamental circuit of e to obtain some signed circuit C' in which $|C' \setminus S| < |C \setminus S|$, and is compatible with \mathcal{O} but not σ . \square

PROOF OF THEOREM 5.5.7 By Corollary 5.5.4, it suffices to provide a bijection between the partial orientations we are counting and the σ' -compatible orientations of M/S for some suitable acyclic circuit signature σ' .

We will work with the following circuit signature σ' of M/S : given a circuit C in M/S , there exists a unique circuit \hat{C} in $C \cup S$ in M and it contains C , set $\sigma'(C) = \sigma(\hat{C})/S$. We claim that σ' is acyclic. Suppose $\sum a_C \sigma'(C) = 0$ for some non-negative a_C 's that are not all zeros, consider $\sum a_C \sigma(\hat{C})$. On one hand the support of the sum is a subset of S , but on the other hand it is an element in the circuit space of M , so it must be zero, which contradicts the acyclicity of σ .

Now we claim that the map $\mathcal{O} \mapsto \mathcal{O}/S$ is the bijection we wanted. By the construction of σ' , such map sends σ -compatible orientations of M to σ' -compatible orientations of M/S . Since every partial orientation we are considering is the same over S , injectivity is trivial. Let \mathcal{O} be a σ' -compatible orientation, we construct a partial orientation $\hat{\mathcal{O}}$ of M as follows: we assign 0 to the elements in S , then orient other elements in \overline{S} according to the orientation of their fundamental circuits specified by σ , and finally assign the orientations of the remaining elements according to \mathcal{O} . By Lemma 5.5.8, if $\hat{\mathcal{O}}$ is not σ -compatible,

then there exists a signed circuit $C \subset S \cup (M \setminus \overline{S})$ that is compatible with $\hat{\mathcal{O}}$ but not σ . By definition, C/S is not compatible with σ' , but it is compatible with \mathcal{O} , a contradiction. Therefore $\hat{\mathcal{O}}$ is σ -compatible and $\hat{\mathcal{O}}/S = \mathcal{O}$, so our map is surjective as well. \square

Remark. Even when σ is of the (highly combinatorial) form described in Example 5.1.4, we do not know a combinatorial proof of Theorem 5.5.7 that does not use the notion of acyclicity, as the constraint on the bi-oriented part seems to make standard deletion-contraction arguments or other similar classical techniques impossible. From a higher level point of view, it is because “lexicographic data” is not closed under arbitrary minor operations (but the notion of acyclicity is).

Now we deduce Theorem 5.5.2 from Theorem 5.5.7, the argument below was due to Backman, who posed Theorem 5.5.7 as a conjecture.

PROOF OF THEOREM 5.5.2: Every continuous orientation \mathcal{O} of M can be viewed as a partial orientation $\hat{\mathcal{O}}$ by declaring $\hat{\mathcal{O}}(e) = \text{sign}(\mathcal{O}(e))$ if $\mathcal{O}(e) = \pm 1$, and $\hat{\mathcal{O}}(e) = 0$ otherwise. Fix an acyclic circuit signature σ . By Theorem 5.3.4 and Proposition 5.3.5, every face of the zonotopal tiling induced by σ corresponds to a partial orientation, namely the combinatorial type of any σ -compatible continuous orientation in the relative interior of the face. Moreover, the dimension of the face equals the size of the bi-oriented part.

Every (d -dimensional) face of the zonotopal tiling is a unimodular lattice parallelotope, hence its q -fold dilation contains $(q - 1)^d$ many interior lattice points. Summing over all faces and using the aforementioned correspondence between faces and σ -compatible partial orientations, we have $E_{Z_M}(q) = \sum_{\mathcal{O}} (q - 1)^{d(\mathcal{O})}$, where the sum is taken over all σ -compatible partial orientations and $d(\mathcal{O})$ is the size of the bi-oriented part of \mathcal{O} . Applying Theorem 5.5.7, we further have $E_{Z_M}(q) = \sum_S T_{M/S}(2, 1)(q - 1)^{|S|}$, where the sum is taken over all independent sets.

By a theorem of Kung [77], we have the following convolution formula:

$$T_M(\lambda\xi + 1, xy + 1) = \sum_{S \subset M} \lambda^{r(M)-r(S)} (-y)^{|S|-r(S)} T_{M[S]}(1 - \lambda, 1 - x) T_{M/S}(1 + \xi, 1 + y).$$

Substituting $(x, y, \lambda, \xi) = (1, 0, \frac{1}{q}, 1)$, we have

$$T_M(1 + \frac{1}{q}, 1) = \sum_{S \subset M} q^{r(S)-r(M)} 0^{|S|-r(S)} T_{M[S]}(1 - \frac{1}{q}, 0) T_{M/S}(2, 1).$$

A term on the right hand side is non-zero only if $|S| = r(S)$, that is, S is independent. But in such case, $q^{r(S)} T_{M[S]}(1 - \frac{1}{q}, 0) = q^{|S|} (1 - \frac{1}{q})^{|S|} = (q - 1)^{|S|}$. Hence the right hand side of the equation is $q^{-r} \sum_S T_{M/S}(2, 1) (q - 1)^{|S|}$ (sum over independent S 's) as desired. \square

Remark. By applying Kung's convolution formula to Stanley's formula, we can deduce that for any d , the sum of $T_{M/S}(2, 1)$'s over all independent sets S of size d equals the number of σ -compatible partial orientations whose bi-oriented part is of size d . But *a priori* we cannot isolate equalities for individual S 's.

5.6 Group Action–Tiling Duality

We establish a connection between group actions of the Jacobian of a regular matroid on the set of bases induced by geometric bijections, and tilings of the cocircuit space that extend zonotopal tilings.

Proposition 5.6.1 *Let A be a totally unimodular matrix of full row rank. Then the zonotope \widetilde{Z}_A equals, up to translation, a full-dimensional cell of the Voronoi decomposition of the row space V^* of A [108] with respect to the cocircuit lattice of $M(A)$. In particular, \widetilde{Z}_A tiles the row space of A in a facet to facet manner, with the period being the cocircuit lattice of A .*

PROOF: By [107, Lemma 3.2 and 3.3], the (affine span of) facets of the Voronoi cell

containing the origin are $\{\mathbf{x} : \mathbf{x} \cdot K = \frac{1}{2}K \cdot K\}$'s for each cocircuit K (as a $\{1, 0, -1\}$ -vector in the cocircuit lattice) of A . By Proposition 5.3.7 and the fact that \mathbf{v}_K there coincides with K when A is totally unimodular, the (affine span of) facets of \widetilde{Z}_A are $\{\mathbf{x} : \mathbf{x} \cdot K = \frac{1}{2}(K + \sum_{e \in E} \pi_{V^*}(\mathbf{e})) \cdot K\}$'s. Therefore \widetilde{Z}_A and the aforementioned Voronoi cell differ only by a translation of $\frac{1}{2} \sum_{e \in E} \pi_{V^*}(\mathbf{e})$. \square

In particular, if we start with a unimodular zonotope $Z := Z_M$ with a chosen zonotopal tiling, then such tiling pattern can be extended to $V^*(M)$ by tiling the whole space with Z . We call such tiling of $V^*(M)$ the *extension* of the zonotopal tiling (of Z).

For the next lemma, recall that by Section 4.2, there is a canonical group action \cdot of $\text{Jac}(M)$ on $\mathcal{G}(M)$. For the sake of clarity, we write the group action as addition, i.e., $g + X := g \cdot X$ for $g \in \text{Jac}(M)$ and $X \in \mathcal{G}(M)$. Moreover, as the group action is simply transitive, for $X, Y \in \mathcal{G}(M)$, it is well-defined to denote by $X - Y$ the unique group element g such that $g + Y = X$.

Lemma 5.6.2 *Let M be a regular matroid, and let $\beta_1, \beta_2 : \mathcal{B}(M) \rightarrow \mathcal{G}(M)$ be two bijections. Then the two $\text{Jac}(M)$ -group actions $g \cdot_1 B := \beta_1^{-1}(g + \beta_1(B))$ and $g \cdot_2 B := \beta_2^{-1}(g + \beta_2(B))$ are isomorphic if and only if there exists $g_0 \in \text{Jac}(M)$ such that $\beta_1(B) - \beta_2(B) = g_0$ for every $B \in \mathcal{B}(M)$.*

PROOF: “ \Rightarrow ”. Pick an arbitrary $B_0 \in \mathcal{B}(M)$ and denote by g_0 the difference $\beta_1(B_0) - \beta_2(B_0)$. For any other $B \in \mathcal{B}(M)$, let $g := \beta_2(B_0) - \beta_2(B)$. We have $g \cdot_1 B = g \cdot_2 B = \beta_2^{-1}((\beta_2(B_0) - \beta_2(B)) + \beta_2(B)) = B_0$. Applying β_1 to both sides yield $g + \beta_1(B) = (\beta_2(B_0) - \beta_2(B)) + \beta_1(B)$ and $\beta_1(B_0) = \beta_2(B_0) + g_0$, respectively. Comparing gives $\beta_1(B) - \beta_2(B) = g_0$.

“ \Leftarrow ”. For any $g \in \text{Jac}(M)$ and $B \in \mathcal{B}(M)$, $g \cdot_1 B = \beta_1^{-1}(g + \beta_1(B)) = \beta_1^{-1}(g + g_0 + \beta_2(B)) = \beta_1^{-1}(g_0 + \beta_2(g \cdot_2 B)) = \beta_1^{-1}(\beta_1(g \cdot_2 B)) = g \cdot_2 B$. \square

Theorem 5.6.3 *Let M be a regular matroid. Let σ be an acyclic circuit signature and let τ, τ' be two acyclic cocircuit signatures. Then $\beta_{\sigma, \tau}, \beta_{\sigma, \tau'}$ induce the same $\text{Jac}(M)$ -group action on $\mathcal{B}(M)$ up to isomorphism if and only if the two extensions of the zonotopal tilings of Z_{M^*} induced by τ, τ' , respectively, are equivalent up to translation.*

PROOF: “ \Rightarrow ”. Suppose $\beta_{\sigma, \tau}, \beta_{\sigma, \tau'}$ induce isomorphic group actions. For each basis B of M , denote by $\gamma_B \in C_1(M)$ the difference $\hat{\beta}_{\sigma, \tau}(B) - \hat{\beta}_{\sigma, \tau'}(B)$, that is, γ_B is the sum of (signed) elements whose orientations differ between $\hat{\beta}_{\sigma, \tau}(B)$ and $\hat{\beta}_{\sigma, \tau'}(B)$. By Lemma 5.6.2, $[\gamma_B] = [\gamma_{B'}] \in \text{Jac}(M)$ for any pair of bases B, B' , which means $\gamma_B - \gamma_{B'} \in Z_1(M) \oplus B_1(M)$.

Fix an arbitrary basis B_0 of M . We claim that in the extension of the zonotopal tiling of Z_{M^*} induced by τ , there is a copy of Z_{M^*} that is $\pi_V(\gamma_{B_0})$ away from the reference zonotope, and its tiling pattern is equal to the zonotopal tiling induced by τ' . Here π_V is the projection onto the cocircuit space of M^* , which is also the circuit space of M .

Let B be any basis of M and bi-orient the elements of its dual basis $B^* = E \setminus B$ in an arbitrary way. Consider the τ -compatible (continuous) orientation \mathcal{O} and τ' -compatible (continuous) orientation \mathcal{O}' extending the bi-oriented B^* . The difference between $\mathcal{O}, \mathcal{O}'$ is precisely γ_B by Proposition 5.3.5, hence the points representing $\mathcal{O}, \mathcal{O}'$ via ψ (and L) in $\widetilde{Z_M}$ differ by a displacement of $\pi_V(\gamma_B)$. In particular, the copies of $\widetilde{Z(B^*)}$ in the two zonotopal tilings of a fixed copy of Z_{M^*} , induced by τ and τ' , respectively, also differ by the displacement of $\pi_V(\gamma_B)$.

The projection map π_V is the zero map over $B_1(M)$ and the identity map over $Z_1(M)$. So the observation that $\gamma_B - \gamma_{B_0} \in Z_1(M) \oplus B_1(M)$ implies $\pi_V(\gamma_{B_0}) \in \pi_V(\gamma_B) + Z_1(M)$ for any B . Since $Z_1(M)$ is the period of any extension tiling, in the comparison of the positions of cells in two zonotopal tilings above, the displacement $\pi_V(\gamma_B)$ can be replaced uniformly by $\pi_V(\gamma_{B_0})$ if we work with their extensions (and consider cells from the translated copies of Z_{M^*} if necessary), and this is precisely our claim.

“ \Leftarrow ”. The proof is essentially the converse of the above argument. In the zonotopal

tilings induced by τ, τ' , respectively, the cells correspond to the same basis B differ by the displacement of $\pi_V(\gamma_B)$, where $\gamma_B = \hat{\beta}_{\sigma, \tau}(B) - \hat{\beta}_{\sigma, \tau'}(B) \in C_1(M)$ is the same element as above. The translation invariance implies that $z_{B, B'} := \pi_V(\gamma_B) - \pi_V(\gamma_{B'}) \in Z_1(M)$ for any bases B, B' . Now $\gamma_B - \gamma_{B'} - z_{B, B'} \in V^*(M)$ as the projection of such element onto $V(M)$ is 0, but $C_1(M) \cap V^*(M) = B_1(M)$ by Proposition 2.3.16, so $\gamma_B - \gamma_{B'} \in Z_1(M) \oplus B_1(M)$. It follows from Lemma 5.6.2 that the group actions induced by $\beta_{\sigma, \tau}, \beta_{\sigma, \tau'}$ are isomorphic. \square

5.7 The ABKS Decomposition as a Zonotopal Tiling

We explain how the ABKS decomposition of a graph G is related to zonotopal tilings. In particular, we describe a family of bijections between spanning trees and (integral) break divisors of a graph that comes from the geometric bijections defined in previous sections. For that reason, we call these bijections as geometric bijections (coming from the ABKS decomposition) as well.

Let Γ be the metric realization of G whose edge lengths are 1, fix a vertex q of G as well as a reference orientation $\overrightarrow{u_i v_i}$ for each edge e_i of G . Let $D = (c_1) + \dots + (c_g)$ be a break divisor. Pick a spanning tree T that induces D , that is, c_i is a chip inside edge $e_i \notin T$ (by a permutation of edge indexes if necessary) and is of distance α_i from u_i . We associate a continuous orientation \mathcal{O}_D of $M(G)^*$ to D : for an element $e \in T$, orient e away from q with respect to T , and set $\mathcal{O}_D(e_i) = 2\alpha_i - 1$ for $i = 1, 2, \dots, g$.

Lemma 5.7.1 *For a break divisor D , \mathcal{O}_D is well-defined up to continuous cocircuit reversals with respect to $M(G)^*$, and they are all σ_q -compatible. Here we interpret σ_q , the q -connected signature defined in Example 5.1.5, as an acyclic circuit signature of $M(G)^*$.*

PROOF: Let \mathcal{O}_D and \mathcal{O}'_D be two continuous orientations constructed from D using different choices of T and e_i 's. When a chip c is in the interior of an edge e , there would be no ambiguity and we must have $\mathcal{O}_D(e) = \mathcal{O}'_D(e)$. If we remove these chips that are in the

interior of edges from D to obtain D' , and remove the corresponding set S of edges from G to obtain G' , then D' is still a break divisor of G' and it is integral. Now $\mathcal{O}_D|_{E \setminus S}, \mathcal{O}'_D|_{E \setminus S}$ are discrete orientations that induce the same orientable divisor $D' - (q)$, so they differ by a series of cycle reversals in G' by Proposition 4.1.5, which are cocircuit reversals in $M(G)^*/S$. The first half of the lemma follows from the observation that these cocircuit reversals can be lifted back to $M(G)^*$ to transform \mathcal{O}_D into \mathcal{O}'_D without involving elements in S .

By the explanation in Example 5.1.5, the orientation of elements in T can be thought as coming from the greedy procedure in the proof of Proposition 5.3.5 with respect to σ_q and $E \setminus T \in \mathcal{B}(M(G)^*)$. By the same proposition, the overall orientation is σ_q -compatible regardless the orientation outside T . \square

The ABKS decomposition can be lifted to the universal cover $H_1(G; \mathbb{R})$ of the tropical Jacobian, while by Proposition 5.6.1, a zonotopal tiling of $\widetilde{Z}_{M(G)^*}$ can be extended to a tiling of $V^*(M(G)^*) \cong V(M(G)) \cong H_1(G; \mathbb{R})$ by tiling the space using $\widetilde{Z}_{M(G)^*}$. Moreover, Proposition 5.6.1 implies that in such a tiling, a facet of $\widetilde{Z}_{M(G)^*}$ will overlap with the opposite facet of another copy of $\widetilde{Z}_{M(G)^*}$.

Theorem 5.7.2 *Let q be an arbitrary vertex of G . Then the following two tilings of $H_1(G; \mathbb{R})$ coincide up to translation: the ABKS decomposition lifted to its universal cover, and the zonotopal tiling of $\widetilde{Z}_{M(G)^*}$ induced by the q -connected signature σ_q of $M(G)^*$, extended to the whole affine span.*

PROOF: Both the ABKS decomposition and the zonotopal tiling induced by σ_q have the same set of cells, namely, one parallelotope \mathcal{C}_T (resp. $\widetilde{Z}(E \setminus T)$) for each spanning tree (resp. basis) T , both equal $\prod_{e \notin T} \pi_{H_1(G; \mathbb{R})}(e)$ up to translation. So it suffices to show that the cells in the ABKS decomposition glue to each other according to the rule described in Proposition 5.3.7.

Let \mathcal{C}_T be the cell in the ABKS decomposition corresponding to a spanning tree T , and

let F be a facet of \mathcal{C}_T corresponding to break divisors whose chip in $e = uv \notin T$ is at v . Pick a break divisor D contained in the relative interior of F (so all chips of D , except the one at v , are in the interior of edges). Let C be the fundamental cycle of e with respect to T and let $f = u'v$ be the other edge in C incident to v . Move the chip at v in D to the interior of f to obtain a new divisor D' , then D' is a break divisor associated to $T' := T \setminus \{e\} \cup \{f\}$, and thus in $\mathcal{C}_{T'}$.

Case I: $v \neq q$. If we choose T to be the spanning tree in the construction of \mathcal{O}_D , then f is oriented from u' to v , and $\mathcal{O}_{D'}$ can be chosen as the continuous orientation obtained from bi-orienting f in \mathcal{O}_D . Since D' is break, $\mathcal{O}_{D'}$ is σ_q -compatible and the edge f we have chosen here is the same “ f ” in Case 2 of the statement of Proposition 5.3.7, so $Z(\widetilde{E \setminus T}) \cong \mathcal{C}_T$ shares the common face F with $Z(\widetilde{E \setminus T'}) \cong \mathcal{C}_{T'}$.

Case II: $v = q$. If we again choose T as the spanning tree to construct \mathcal{O}_D , then f is oriented from $v = q$ to u' and C is a directed cycle in \mathcal{O}_D . If we choose T' instead to construct another continuous orientation \mathcal{O}'_D corresponding to D , then \mathcal{O}_D and \mathcal{O}'_D differ by a reversal of C , and \mathcal{O}'_D is on the opposite facet F' of F in $\widetilde{Z_{M(G)^*}}$. Now $\mathcal{O}_{D'}$ can be chosen as the continuous orientation obtained from bi-orienting f in \mathcal{O}'_D , so F' is a facet of $Z(\widetilde{E \setminus T'})$. In particular, in the extension tiling, $Z(\widetilde{E \setminus T})$ in a copy of $\widetilde{Z_{M(G)^*}}$ shares the common face F with $Z(\widetilde{E \setminus T'})$ in other copy of $\widetilde{Z_{M(G)^*}}$, which is Case 1 of the statement of Proposition 5.3.7. \square

Since the shifting part of geometric bijections only relies on the local structure of cells, Theorem 5.7.2 yields the following version of “geometric bijections”.

Corollary 5.7.3 *Let σ be an acyclic cycle signature of G . Given a spanning tree T of G , let D_T be the divisor of G in which we orient each $e \notin T$ according to its orientation in $\sigma(C(T, e))$ and then put a chip at the head of e . Then the map $\beta_\sigma : T \mapsto D_T$ is a bijection from $\mathcal{ST}(G)$ to $\mathcal{BD}(G)$.*

In view of Theorem 5.6.3, we may ask what is the “dual” group action of the ABKS

decomposition. An answer is provided in Section 6.2.1.

5.8 Algorithmic Aspects

5.8.1 Inverse Algorithm for Geometric Bijections via Linear Programming

We describe an inverse algorithm which furnishes an inverse to the map ϕ , and hence to $\hat{\beta}$ and β . The key ingredient is that given a continuous orientation \mathcal{O} and a vector \mathbf{w} that induces acyclic signatures σ and σ^* with respect to a matrix A (cf. Lemma 5.3.2), the unique σ -compatible continuous orientation \mathcal{O}' in the continuous circuit reversal class of \mathcal{O} (whose existence is guaranteed by Theorem 5.3.4) can be computed in polynomial-time using linear programming.

To do so, we solve the “max-flow” linear program

$$\max\{\mathbf{w}^T \mathbf{y} : A\mathbf{y} = 0, -1 - \mathcal{O}(e) \leq \mathbf{y}_e \leq 1 - \mathcal{O}(e) \forall e\}. \quad (5.2)$$

Let $\tilde{\mathbf{y}}$ be an optimal solution. Consider the continuous orientation \mathcal{O}' given by $\mathcal{O}'(e) = \mathcal{O}(e) + \tilde{\mathbf{y}}_e$ for every e , we claim that this is the σ -compatible continuous orientation we are looking for. The condition $-1 - \mathcal{O}(e) \leq \mathbf{y}_e \leq 1 - \mathcal{O}(e)$ in the linear program guarantees that \mathcal{O}' is a valid continuous orientation, and the condition $A\mathbf{y} = 0$ guarantees that \mathcal{O}' is circuit-reversal equivalent to \mathcal{O} . Finally, the orientation \mathcal{O}' is σ -compatible: indeed, if \mathcal{O}' is compatible with some $-\sigma(C)$, then one can easily check that $\tilde{\mathbf{y}} + \delta \mathbf{v}_{\sigma(C)}$ is also a feasible solution to the linear program for sufficiently small $\delta > 0$, contradicting the optimality of $\tilde{\mathbf{y}}$.

The linear program (5.2) and its dual version, together with the fact that one can work independently in the circuit (resp. cocircuit) part of an orientation (Lemma 4.1.3), imply the following.

Proposition 5.8.1 *There is a polynomial-time algorithm to compute the unique (σ, σ^*) -compatible continuous orientation equivalent to a given continuous orientation.*

Now suppose we are given a (σ, σ^*) -compatible discrete orientation $\mathcal{O} \in \chi(M; \sigma, \sigma^*)$ in the general realizable matroid case, or any discrete orientation in the regular case. We will work with the general case, as the latter case can be reduced to it using Proposition 5.8.1. Suppose \mathcal{O} was shifted into the cell $\widetilde{Z}(\widetilde{B})$ after moving by a displacement of $-\pi_{V^*}(\mathbf{w})$ (for sufficiently short \mathbf{w}). By solving the linear program (5.1) in the proof of Theorem 5.4.2, we obtain a vector \mathbf{f} such that a continuous orientation in the cell $Z(B)$ is given by $\widetilde{\mathcal{O}} := \mathcal{O} - 2\mathbf{f}$. Then we apply Proposition 5.8.1 to obtain the (σ, σ^*) -compatible continuous orientation \mathcal{O}' equivalent to $\widetilde{\mathcal{O}}$, the basis B will then be the set of bi-oriented elements in \mathcal{O}' .

Since linear programming admits a polynomial-time algorithm [98, Chapter 13], we only need to guarantee the bit complexity is of polynomial-size with respect to the input through out the algorithm. But precisely because linear programming can be solved in polynomial-time, the output of any intermediate linear program has polynomial bit complexity with respect to the input. Finally, \mathbf{w} can be taken as $(1, 1/2, 1/4, \dots, 1/2^{m-1})$ (Example 5.1.4), which has polynomial bit size with respect to A .

Summarizing the discussion, we have the following theorem.

Theorem 5.8.2 *There is a polynomial-time algorithm to compute the inverse of $\hat{\beta}$ and β .*

5.8.2 Sampling Algorithms

By mimicking the strategy from [12], we can now produce a polynomial-time algorithm for randomly sampling bases of a regular matroid, which gives an answer to the question posed by Jeremy Martin and Farbod Shokrieh at the AIM chip-firing workshop [67]. The high-level strategy is:

1. Fix an arbitrary reference orientation \mathcal{O} of M , as well as an acyclic circuit signature σ and an acyclic cocircuit signature σ^* .
2. Compute $\text{Jac}(M) \cong \text{coker}(AA^T)$ as a product of cyclic groups by computing the Smith normal form of AA^T [72].
3. Use this presentation to choose a random element $\text{Jac}(M)$, and transform it to an element $[\gamma] \in \text{Jac}(M) \cong \frac{C_1(M)}{Z_1(M) \oplus B_1(M)}$.
4. Compute $[\mathcal{O}'] := [\gamma] \cdot [\mathcal{O}] \in \mathcal{G}(M)$, where \cdot is the polynomial-time computable group action from Section 4.2.
5. Compute the (σ, σ^*) -compatible orientation \mathcal{O}'' in $[\mathcal{O}']$, which can be done in polynomial-time by Proposition 5.8.1.
6. Output the set of bi-oriented elements in \mathcal{O}'' as the random basis.

Besides theoretical interests, a random basis sampling algorithm for graphic matroids (that is, a random spanning tree sampling algorithm) has real-life applications [52]. And while the current implementation of our algorithm is slower than some other available algorithms (e.g. the random walk based ones such as [111], or the determinant based ones such as [45]), a feature of our algorithm is that it uses information theoretically minimum amount of randomness, namely, it suffices to use $\log_2 |\mathcal{B}(M)|$ random bits.

Using the combinatorial description of lattice points in a (unimodular) zonotope (cf. Section 5.5), we can deduce another algorithmic result related to sampling. The combinatorial part was first studied by Tetali et al. [17], in which they posed the problem of finding a polynomial-time algorithm for sampling acyclic orientations of G with a unique source q .

Proposition 5.8.3 *Let G be a connected graph without loops. Then for any vertex q , the number of acyclic orientations of G with a unique source q is equal to $T_G(1, 0) = T_{M(G)^*}(0, 1)$.*

Furthermore, the problem of uniformly sampling an acyclic orientation with a unique source is polynomial-time equivalent to uniformly sampling an interior lattice point of $Z_{M(G)^*}$.

PROOF: By Theorem 5.5.3, the interior lattice points of $Z_{M(G)^*}$ are precisely the acyclic cut reversal classes of G . We then claim that their σ_q -compatible representatives are exactly the acyclic orientations of G with a unique source q , from which the enumerative statement follows from Theorem 5.5.3 immediately. If an acyclic orientation \mathcal{O} contains another source vertex q' , then $\delta(q')$ must be pointing away from q' thus pointing towards q , hence \mathcal{O} is not σ_q -compatible. Conversely, if q is the unique source of an acyclic orientation \mathcal{O} and K is a directed cut pointing from U to $V(G) \setminus U$ such that $q \notin U$, then $\mathcal{O}|_U$ must contain a source $q' \neq q$, which is also a source in \mathcal{O} .

Using the linear algebraic map ψ , we have an easy correspondence between the interior lattice points of $Z_{M(G)^*}$ and the acyclic cut reversal classes of G ; and given any acyclic orientation of G , the unique σ_q -compatible orientation equivalent to it can be computed in polynomial-time either using Proposition 5.8.1 or the greedy procedure described in Proposition 4.3.8 (which runs in polynomial-time by the remark followed). This establishes the equivalence. \square

Despite we cannot settle the sampling problem of Tetali et al. at this moment, the reduction shed some light to the problem as there are works on sampling lattice points of a polytope using tools from high dimensional geometry [73], so advances along these directions might help us to solve the original problem; conversely, combinatorial hardness results can in turn be translated into hardness results of geometric problems. As another note, in [17], the authors also gave explicit and polynomial-time computable bijections between acyclic orientations with a unique source and other combinatorial objects, such as *spanning trees with zero external activity* (as known as “safe tree” [75]) and *maximal G -parking functions*, so our reduction can be applied to the corresponding sampling problems as well.

5.8.3 Edge Ordering Maps and Their Combinatorial Inverses

We consider a special class of geometric bijections from the ABKS decomposition. An advantage of these bijections is that they admit a purely combinatorial polynomial-time inverse algorithm.

Fix a total ordering and a reference orientation of the edges $e_1 = \overrightarrow{u_1v_1} < e_2 = \overrightarrow{u_2v_2} < \dots < e_m = \overrightarrow{u_mv_m}$ of G . By Example 5.1.4, such data induces an acyclic circuit signature. Hence the following *edge ordering map* is a bijection between spanning trees and break divisors of G .

Input: A spanning tree T of G .
Output: A break divisor $D \in \text{Div}(G)$.
Set $D := 0$.
for $f \notin T$ **do**
 $C :=$ Unique cycle contained in $T + f$
 $i :=$ index of the smallest edge in C
 Orient f to have the same orientation as e_i in C
 $D := D + \text{Head}(f)$
end
Output D .

Algorithm 2: Edge Ordering Map

Now we give a combinatorial inverse algorithm to Algorithm 2, thereby providing a combinatorial proof that an edge ordering map is indeed a bijection.

We will first describe the key subroutine *Inverse* in Algorithm 3, which works at the level of orientations; then we will give the main algorithm in Algorithm 4. Here the subroutine *DivisorToOrientation* is the algorithm by Backman [4, Algorithm 7.6] which, given a break divisor D on a graph G and a vertex q , outputs a q -connected orientation \mathcal{O} such that $D = D_{\mathcal{O}} + (q)$.

With the routine *Inverse*, the overall inverse algorithm, Algorithm 4, is fairly simple.

Theorem 5.8.4 *Algorithm 4 always terminates and it is the inverse of Algorithm 2.*

Input: A connected graph H , a vertex $q \in V(H)$ and a q -connected orientation \mathcal{O} .

Output: A spanning tree T of H .

if H is a tree **then**

 | $T := H$

else

 | $i :=$ index of the smallest edge in H

 | $P :=$ A directed path from q to v_i in \mathcal{O}

 | Reverse the edges of P in \mathcal{O} to obtain $\hat{\mathcal{O}}$

 | $U :=$ Vertices u_i can reach in $\hat{\mathcal{O}}$

 | **if** $v_i \in U$ **then**

 | **if** e_i goes from v_i to u_i in $\hat{\mathcal{O}}$ **then**

 | $Q :=$ A directed path from u_i to v_i in $\hat{\mathcal{O}}$

 | Reverse the edges of the directed cycle $\{e_i\} \cup Q$ in $\hat{\mathcal{O}}$

 | **end**

 | $T := \text{Inverse}(H - e_i, v_i, \hat{\mathcal{O}} - e_i)$

 | **else**

 | $T := \{e_i\} \cup \text{Inverse}(H[U], u_i, \hat{\mathcal{O}}|_{H[U]}) \cup \text{Inverse}(H[U^c], v_i, \hat{\mathcal{O}}|_{H[U^c]})$

 | **end**

end

Output T .

Algorithm 3: The subroutine *Inverse*

Input: A connected graph G and a break divisor $D \in \text{Div}(G)$.

Output: A spanning tree T of G .

$\mathcal{O} := \text{DivisorToOrientation}(G, D, v_1)$

$T := \text{Inverse}(G, v_1, \mathcal{O})$

Output T .

Algorithm 4: Inverse to the Edge Order Map

PROOF: We shall first prove that every recursive step in *Inverse* is valid, then show by induction on the number of edges that *Inverse* is the inverse of Algorithm 2, namely, if $T = \text{Inverse}(H, q, \mathcal{O})$, then the break divisor associated to T via Algorithm 2 is $D_{\mathcal{O}} + (q)$.

The case of H being a tree is obviously correct. For the non-trivial cases, first notice that $D_{\hat{\mathcal{O}}} + (v_i) = D_{\mathcal{O}} + (q)$, so they correspond to the same break divisor. Furthermore, $\hat{\mathcal{O}}$ is v_i -connected: for any vertex u in H , pick a directed path R from q to u in \mathcal{O} , let w be the vertex in the intersection of $V(P)$ and $V(R)$ that is closest to v_i on P (w could be q itself), then concatenating the portion of P from v_i to w and the portion of R from w to u gives a directed path from v_i to u in $\hat{\mathcal{O}}$.

In the case of $v_i \in U$, the algorithm performs a cycle reversal if necessary to guarantee that e_i goes from u_i to v_i in $\hat{\mathcal{O}}$, so $\hat{\mathcal{O}} - e_i$ is a v_i -connected orientation of $H - e_i$ as v_i never needed to use e_i to reach any other vertices in the new $\hat{\mathcal{O}}$. Hence the recursive call $\text{Inverse}(H - e_i, v_i, \hat{\mathcal{O}} - e_i)$ is valid. Letting $T = \text{Inverse}(H - e_i, v_i, \hat{\mathcal{O}} - e_i)$, by induction the break divisor obtained from T in $H - e_i$ using Algorithm 2 is $D_{\hat{\mathcal{O}} - e_i} + (v_i)$. Consider the fundamental cycle C of e_i in $T + e_i \subset H$. Since e_i is the smallest edge in C , e_i will be oriented as its own reference orientation $\overrightarrow{u_i v_i}$ in Algorithm 2, hence the break divisor obtained from T in H using Algorithm 2 is indeed $D_{\hat{\mathcal{O}} - e_i} + 2(v_i) = D_{\hat{\mathcal{O}}} + (v_i)$.

In the last case with $v_i \notin U$, every edge between U and U^c goes from U^c to U in $\hat{\mathcal{O}}$. In particular, $D_{\hat{\mathcal{O}}} = D_{\hat{\mathcal{O}}|_{H[U]}} + D_{\hat{\mathcal{O}}|_{H[U^c]}} + \sum_{e=uv', u \in U, v' \in U^c} (u)$, here we abuse notations and consider $D_{\hat{\mathcal{O}}|_{H[U]}}$ and $D_{\hat{\mathcal{O}}|_{H[U^c]}}$ as divisors on H in the obvious way. On one side, $\hat{\mathcal{O}}$ restricted to $H[U^c]$ is v_i -connected as v_i could never use edges between U and U^c to access vertices in U^c , hence $\text{Inverse}(H[U^c], v_i, \hat{\mathcal{O}}|_{H[U^c]})$ is a valid call; on the other side, $\hat{\mathcal{O}}$ restricted to $H[U]$ is u_i -connected by the construction of U , hence $\text{Inverse}(H[U], u_i, \hat{\mathcal{O}}|_{H[U]})$ is also a valid call. Suppose T is the outputted tree here, we consider the break divisor associated to T via Algorithm 2. By induction hypothesis, the contribution of those non-tree edges in $H[U]$ and $H[U^c]$ is equal to $D_{\hat{\mathcal{O}}|_{H[U]}} + (u_i) + D_{\hat{\mathcal{O}}|_{H[U^c]}} + (v_i)$. For every edge $f \neq e_i$ between U and U^c , the fundamental cycle of $T + f$ contains e_i , so f will be oriented

from U^c to U , thus these edges contribute $[\sum_{e=uu', u \in U, u' \in U^c} (u)] - (u_i)$ to the final divisor. Summing these contributions of non-tree edges gives $D_{\partial} + (v_i)$ as claimed. \square

Now we analyze the complexity of Algorithm 4. The complexity of the subroutine *DivisorToOrientation* is essentially the complexity of a maximum flow algorithm (here any exact algorithm for unit capacity directed graphs suffices, say the $\tilde{O}(m^{10/7})$ algorithm by Madry [84]). Each instance of the first case of *Inverse* takes $O(|V(H)|)$ time, but any two H 's considered in the computational process of Algorithm 4 are disjoint, so the first case takes $O(n)$ time in total. Lastly, for each non-trivial call of *Inverse*, a different smallest edge e_i is being considered, so there can be $O(m)$ such calls, and it is easy to see such a call can be handled in $O(m)$ time (P , U , and Q can all be found by BFS, and other maintenance/modification operations also take $O(m)$ time). Therefore the total time complexity is $O(m^2)$.

5.9 A Note on Lawrence Polytopes and Lawrence Ideals

In this section, we provide a brief survey on the *Lawrence construction* in combinatorial commutative algebra and polyhedral geometry, and describe how the notion of geometric bijections is related to it. We will assume basic definitions in these areas, and refer the reader to standard texts such as [104] and [115].

Let $\pi : P \rightarrow Q$ be a projection of polytopes and let \mathbf{w} be a vector in the affine span \mathbb{R}^k of P . The π -coherent subdivision with respect to \mathbf{w} is a subdivision of Q consisting of the projection of the “lower faces” of $\{(\pi(\mathbf{x}), \mathbf{w} \cdot \mathbf{x}) : \mathbf{x} \in P\} \in \mathbb{R}^k \times \mathbb{R}$ [115, Section 9.1]. We will be interested in the projection $[0, 1]^E \rightarrow Z$ from a hypercube to a zonotope, in which the *tight* π -coherent subdivisions correspond to zonotopal tilings induced by acyclic circuit signatures, cf. the remark after Theorem 5.3.4. The *fiber polytope* of $\pi : P \rightarrow Q$ is a polytope whose face poset is isomorphic to the poset of all π -coherent subdivisions ordered by refinement [22]; while not a standard terminology, we will call the normal fan of fiber polytope the *fiber fan*.

Let \mathcal{A} be the set of vertices of some polytope P (in general \mathcal{A} can be any point configuration). The *secondary polytope* of P is the fiber polytope of the projection $\pi : \Delta^{|\mathcal{A}|-1} \rightarrow P$ that maps the vertices of the simplex $\Delta^{|\mathcal{A}|-1}$ to the vertices of P [54], the π -coherent subdivisions here are well known as the *regular subdivisions* of P [104, Chapter 8]. The normal fan of the secondary polytope is the *secondary fan* of P .

Let I be a homogeneous ideal in $R = K[x_1, \dots, x_n]$ and let $\mathbf{w} \in \mathbb{R}^n$ be a vector. The monomial term order $<_{\mathbf{w}}$ orders monomials in R by $\mathbf{x}^\alpha <_{\mathbf{w}} \mathbf{x}^\beta$ whenever $\mathbf{w} \cdot \alpha < \mathbf{w} \cdot \beta$. By [104, Theorem 1.11], each monomial initial ideal of I is equal to $\text{in}_{<_{\mathbf{w}}}(I)$ for some \mathbf{w} . By grouping vectors in \mathbb{R}^n according to the initial ideals they induce, we obtain the *Gröbner fan* of I , where its full-dimensional cones correspond to the monomial initial ideals of I by $\mathcal{C} \leftrightarrow \text{in}_{<_{\mathbf{w}}}(I)$ for any $\mathbf{w} \in \mathcal{C}^\circ$. It can be proven that the Gröbner fan is the normal fan of certain polytope known as the *state polytope* [104, Chapter 1–2].

Let A be an $r \times m$ integer matrix. The *Lawrence polytope* of A is the polytope in \mathbb{R}^{r+m} whose vertices are the columns of $\hat{A} = \begin{pmatrix} A & \mathbf{0} \\ \mathbf{I} & \mathbf{I} \end{pmatrix}$, here both identity matrices are of size $m \times m$ [69]. The *Lawrence ideal* J of A is the (homogeneous) ideal in $2m$ variables $x_1, \dots, x_m, y_1, \dots, y_m$ generated by $\mathbf{x}^{\mathbf{v}^+} \mathbf{y}^{\mathbf{v}^-} - \mathbf{x}^{\mathbf{v}^-} \mathbf{y}^{\mathbf{v}^+}$'s for all $\mathbf{v} \in \ker(\hat{A}) \cap \mathbb{Z}^m$ [104, Chapter 7]. Here $\mathbf{v} = \mathbf{v}^+ - \mathbf{v}^-$, $\text{supp}(\mathbf{v}^+) \cap \text{supp}(\mathbf{v}^-) = \emptyset$ is the decomposition of \mathbf{v} into position part and negative part.

When M is a regular oriented matroid, the Lawrence polytope and Lawrence ideal associated to M are the Lawrence polytope and Lawrence ideal of any totally unimodular realization A of M . Notice that in such case, every element $\mathbf{v} \in \ker(\hat{A}) \cap \mathbb{Z}^E$ is of the form $(\mathbf{u}, -\mathbf{u})$ for some $\mathbf{u} \in \ker(A) \cap \mathbb{Z}^E$, and the sign patterns of \mathbf{u} 's are precisely the vectors of M .

We have the following equivalence of fans (and polytopes).

Theorem 5.9.1 *Let M be a regular matroid. Then the following four fans coincide up to lineality spaces.*

1. *The fan in $V^*(M)$ associated to the central hyperplane arrangement consisting of C^\perp 's for each cocircuit C of M .*
2. *The fiber fan of the projection $[0, 1]^E \rightarrow Z_{M^*}$.*
3. *The secondary fan of the Lawrence polytope of M^* .*
4. *The Gröbner fan of the Lawrence ideal J_{M^*} of M^* .*

PROOF: The equivalence of (3) and (4) is [104, Proposition 8.15(a)]. The equivalence of (2) and (3) can be deduced using a special case of the ‘‘Cayley trick’’ [69], which gives a bijection between (coherent) fine zonotopal tilings and (coherent) triangulations of the Lawrence polytope via an explicit polyhedral construction, and in particular it implies the fiber fan in (2) is equal to the secondary fan in (3) up to lineality spaces [103, Theorem 5.1].

Now we prove the equivalence of (4) and (1). By [90, Proposition 7.8], a minimal universal Gröbner basis of the Lawrence ideal is given by $\{\mathbf{x}^{C^+} \mathbf{y}^{C^-} - \mathbf{x}^{C^-} \mathbf{y}^{C^+} : C \in \mathcal{C}^*(M)\}$. Therefore to specify a full-dimensional open cone $\mathcal{C}^\circ \subset \mathbb{R}^{2m}$ in the Gröbner fan of J_{M^*} is the same as to specify the initial term of each basis element with respect to the corresponding monomial term order $\langle_{\hat{\mathbf{w}}}$, $\hat{\mathbf{w}} \in \mathcal{C}^\circ$, which is equivalent to specifying the orientation of each cocircuit of M picked out by $\langle_{\hat{\mathbf{w}}}$. For a cocircuit of M , $\langle_{\hat{\mathbf{w}}}$ picks out the term $\mathbf{x}^{C^+} \mathbf{y}^{C^-}$ if and only if $\mathbf{w} \cdot C > 0$, where $\mathbf{w} \in \mathbb{R}^m$ is the vector given by $\mathbf{w}_i := \hat{\mathbf{w}}_{x_i} - \hat{\mathbf{w}}_{y_i}$. Therefore two generic vectors $\hat{\mathbf{w}}_1, \hat{\mathbf{w}}_2 \in \mathbb{R}^{2m}$ induce different initial ideals if and only if $\pi_{V^*}(\mathbf{w}_1), \pi_{V^*}(\mathbf{w}_2) \in V^*$ are separated by some hyperplane of the form C^\perp , thus the fan in (1) is the Gröbner fan of J_{M^*} quotiented by its lineality space. \square

Remark. The lineality space of the Gröbner fan in (4) is of dimension $2m - r$, m of these dimensions are accounted by the fact that the behavior of $\langle_{\hat{\mathbf{w}}}$ only depends on the differences $\hat{\mathbf{w}}_{x_i} - \hat{\mathbf{w}}_{y_i}$'s, while the fact that the cocircuit signature induced by \mathbf{w} is independent of its circuit part accounts for the remaining $m - r$.

Next we study the relation between the fans in (1) and (2) directly, which can be thought as a special case of the Bohne–Dress theorem. Since the zonotopal tiling of Z_{M^*} induced by \mathbf{w} changes only if the acyclic cocircuit signature of M induced by \mathbf{w} changes, the fan in (1) is a refinement of the fiber fan in (2). Using the setup of geometric bijections, we can give a perhaps more intuitive proof of the fact that they are actually equal. As a corollary, one can replace any step in the above proof by the following argument.

For technical reasons, we will assume the matroid has no loops nor isthmuses. It is not a serious compromise as such elements can be handled separately easily.

Let $\mathbf{w}_1, \mathbf{w}_2 \in V^*(M)$ be two vectors in the interior of two adjacent cones of the fan in (1), separated by the hyperplane C^\perp . Denote by σ_1, σ_2 the cocircuit signatures of M they induce. Fix an arbitrary acyclic circuit signature τ of M , we claim that $\beta_{\tau, \sigma_1}, \beta_{\tau, \sigma_2} : \mathcal{B}(M) \rightarrow \mathcal{G}(M)$ are different bijections. Pick any basis B of M in which C is a fundamental cocircuit respect to it, the cell \mathcal{C}_B associated to B in the zonotopal tiling (of Z_M) induced by τ has two facets parallel to C^\perp . Let \mathcal{O} (as a τ -compatible orientation) be the vertex that get shifted into \mathcal{C}_B along $-\mathbf{w}_1$. From the geometry of parallelotopes, the τ -compatible orientation \mathcal{O}' that get shifted into \mathcal{C}_B along $-\mathbf{w}_2$ differs from \mathcal{O} by precisely one element, which means \mathcal{O} and \mathcal{O}' are not circuit-cocircuit equivalent, in particular $\beta_{\tau, \sigma_1}(B) = [\mathcal{O}] \neq [\mathcal{O}'] = \beta_{\tau, \sigma_2}(B)$. By duality, $\beta_{\sigma_1, \tau}, \beta_{\sigma_2, \tau} : \mathcal{B}(M^*) \rightarrow \mathcal{G}(M^*)$ are different bijections as well. But since the shifting parts are the same, it must be the case that the zonotopal tilings of Z_{M^*} induced by σ_1 and σ_2 are different, and $\mathbf{w}_1, \mathbf{w}_2$ are in different cones in the fiber fan.

Finally, we mention an application of the connections above by giving an algebraic proof of Proposition 5.4.8; the idea and certain special cases were previously studied in [5, Section 4]. Given an acyclic circuit signature σ , there exists a generic monomial term order $<$ corresponding to $-\sigma$ by Theorem 5.9.1. Encode an orientation \mathcal{O} as the monomial $\mathbf{x}^{\mathcal{O}^+} \mathbf{y}^{\mathcal{O}^-}$. Given a signed circuit C , a division by $\mathbf{x}^{C^+} \mathbf{y}^{C^-} - \mathbf{x}^{C^-} \mathbf{y}^{C^+}$ using $\mathbf{x}^{C^+} \mathbf{y}^{C^-}$ as the leading term produces a new remainder if and only if C is compatible with \mathcal{O} , and

such remainder is a monomial representing the orientation obtained from reversing C in \mathcal{O} . From this, we can see that an orientation is σ -compatible if and only if the corresponding monomial is a *standard monomial* with respect to $<$, and Proposition 5.4.8 follows from the existence and uniqueness assertion of remainder of division by a Gröbner basis [36, Section 1.3].

CHAPTER 6

BERNARDI PROCESSES

6.1 Setup of Bernardi Processes

Bernardi introduced a combinatorial process [19, 20] to study the sandpile model on graphs. The key notion used in his work is *ribbon graphs* (also known as combinatorial maps or rotation systems). We will assume all graphs we consider here connected, unless otherwise specified.

Definition 6.1.1 *A ribbon structure on a graph G is a cyclic ordering of edges incident to v for each vertex v of G . A graph together with a ribbon structure is a ribbon graph.*

If G is embedded on an orientable surface, then the orientation of the surface induces a ribbon structure. A non-trivial fact in topological graph theory is that the converse is true.

Theorem 6.1.2 [61, Section 3.2] *The ribbon structures on a graph G are in one-to-one correspondence with the embeddings of G onto some closed orientable surface up to homeomorphism.*

A ribbon graph is *planar* if the surface the graph is embedded onto is a sphere.

Now we describe Bernardi process. Fix a ribbon structure of a graph G and a *starting pair* (v, e) , where e is an edge and v is a vertex incident to e . For any spanning tree T of G , the Bernardi process produces a tour $(v_0, e_0, v_1, \dots, v_k, e_k)$ of the vertices and edges of G . Informally, the tour is obtained by walking along edges belonging to T and cutting through edges not belonging to T , beginning with e and proceeding according to the ribbon structure. Explicitly, start with $v_0 = v, e_0 = e$, and in each step, determine v_{i+1}, e_{i+1} using v_i, e_i as follows: if $e_i \notin T$, then set $v_{i+1} = v_i$ and set e_{i+1} to be the next edge of e_i around

v_i in the cyclic order; otherwise set v_{i+1} to be the other end of e_i and set e_{i+1} to be the next edge of e_i around v_{i+1} . The process stops when every edge is traversed exactly twice.

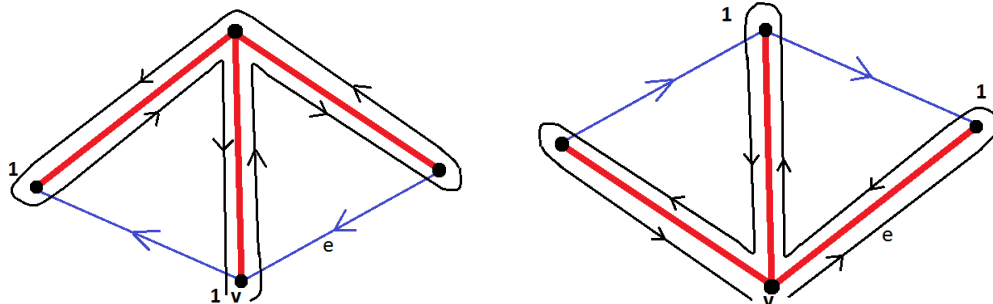


Figure 6.1: Bernardi processes on two different spanning trees, using the same (planar) ribbon structure and starting pair ([13, Figure 2]).

Such tour produces a total ordering $<_T$ (depending on T) of $E(G)$ according to the visiting order by the walk, that is, $f <_T f'$ if the first appearance of f in the tour is earlier than the first appearance of f' . In particular, we can define *Bernardi activities* similar to the classical counterpart using $<_T$'s. An edge $f \notin T$ (resp. $f \in T$) is *externally Bernardi active* (resp. *internally Bernardi active*) if f is the smallest edge in the fundamental cycle (resp. cut) of f with respect to $<_T$. Bernardi gave another description of the Tutte polynomial using his notion of activity.

A Bernardi process produces two natural maps from spanning trees to divisors. Let T be a spanning tree and consider the Bernardi tour obtained from performing the Bernardi process on T . In the first map, we associate a divisor to the tour as follows: each edge not in T appears twice as $e_i, e_j, i < j$ in the tour, for each such edge, add a chip at v_i . The second map is similar to the first, except that we only consider contribution from edges not in T that are not externally Bernardi active, and that for the sake of convenience, we normalize the divisor to degree 0 by modifying the number of chips at v . Bernardi proved the following theorem concerning the two maps.

Theorem 6.1.3 [20] *Given any ribbon graph G and starting edge (v, e) . The first map is a bijection between spanning trees and break divisors, and the second map is a bijection between spanning trees and degree 0 v -reduced divisors.*

In view of Theorem 6.1.3, we call both maps *Bernardi bijections*, and we will specify which one are we referring to upon using the term.

Remark. In the original language of Bernardi, spanning trees are mapped to *v -critical configurations* in the second case, but there is a natural correspondence between critical configurations and reduced divisors [11, Lemma 5.6], so we will use the latter convention.

6.2 Bernardi Bijections as Geometric Bijections

All Bernardi bijections in this section are of the first type. We will prove that every Bernardi bijection of a planar ribbon graph is a geometric bijection coming from the ABKS decomposition, such relation was a conjecture by Baker [13, Remark 5.2]. We will also sketch a partial converse for the non-planar case.

6.2.1 Planar Case

In this section, G will always be a bridgeless planar ribbon graph, in which we implicitly assume it is a plane graph, i.e., a graph already embedded onto \mathbb{R}^2 as a geometric object (cf. [44, Section 4.2]). We first prove that for any Bernardi bijection of G , there exists some cycle signature σ of G so that the bijection is equal to β_σ (cf. Corollary 5.7.3). We call maps of the form β_σ with cycle signature σ (not necessarily acyclic) *cycle orientation maps*.

We adopt the convention that the ribbon structures on plane graphs are induced by the counter-clockwise orientation of the plane. The only exception is when we are working with planar duals, in which case we assume that the ribbon structures of dual graphs are induced by the clockwise orientation of the plane.

Proposition 6.2.1 *The Bernardi bijection with starting pair $(x, e = xy)$ is a cycle orientation map. More explicitly, the cycles are oriented as follows: if e is outside C , orient C clockwise; if e is inside C , orient C counter-clockwise; if e is on C , orient C as the opposite of \overrightarrow{xy} .*

PROOF: Let $C = x_1x_2 \dots x_r$ be a cycle with vertices indexed in counter-clockwise order, and let T be a spanning tree for which C is a fundamental cycle. Without loss of generality, T is missing the edge $\hat{e} := x_r x_1$ from C . Suppose e is outside C . Then the Bernardi tour starting from e , when restricted to C , traverses the “outside” of C in a counter-clockwise manner before going to the “inside” of C by cutting through \hat{e} . Hence the tour will put a chip at x_r , which corresponds to orienting C clockwise. See Figure 6.2 for illustration. The analysis of the remaining cases are similar. \square

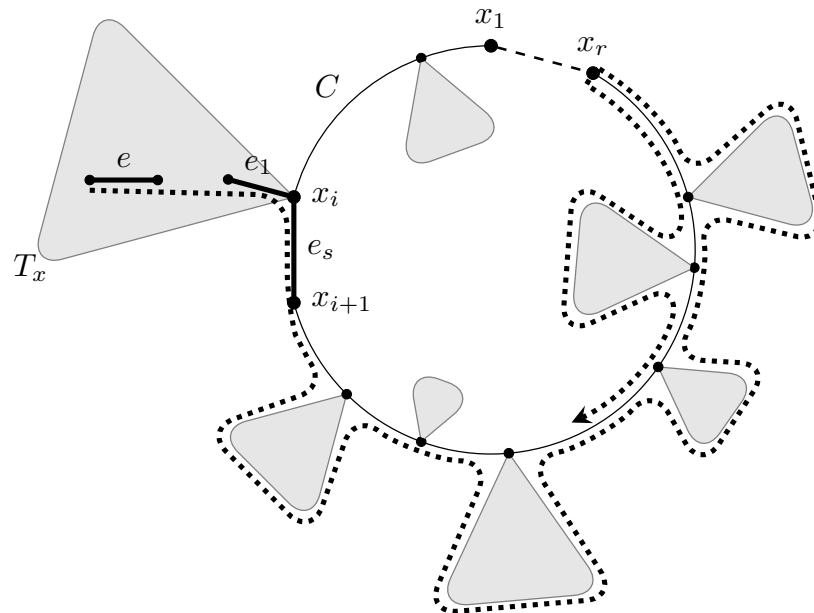


Figure 6.2: Illustration of Proposition 6.2.1. Picture by Emma Cohen.

One immediate observation is that a more correct way to index Bernardi bijections of a plane graph is by faces rather than by starting pairs.

Corollary 6.2.2 *Let $(v, e = uv)$ be a starting pair on G . Let F be the face to the right of (v, e) (cf. Figure 6.3). Then the cycle signature described in Proposition 6.2.1 can be interpreted as follows: a cycle is oriented into counter-clockwise if and only if F is in the interior of the cycle. Conversely every face corresponds to some Bernardi bijection in such manner. Denote by σ_F such cycle signature of G .*

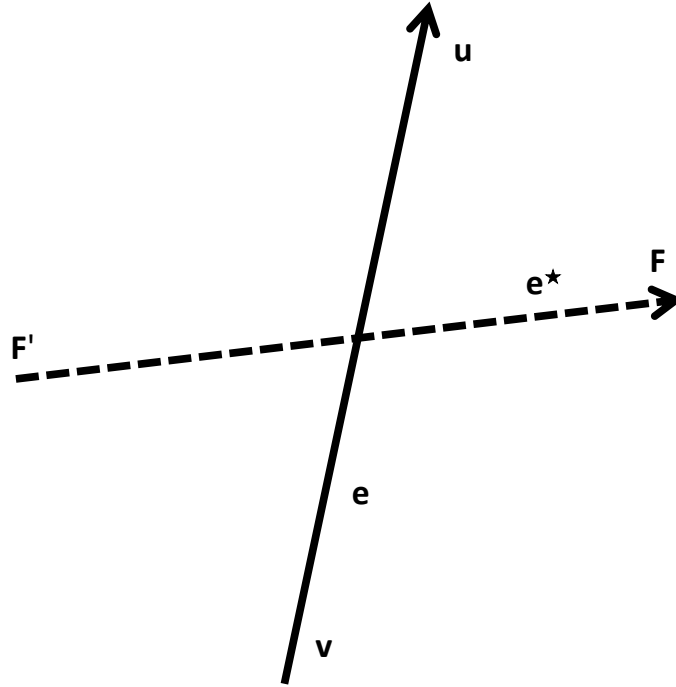


Figure 6.3: Some conventions on the orientation of plane graphs. Here the orientation of e induces the orientation of the dual edge e^* , and F is said to be the face on the right of (v, e) .

Now we prove Baker's conjecture.

Theorem 6.2.3 *Let G^* be the dual graph of G . Let F be a face of G and denote by F^* the corresponding dual vertex of G^* . Then σ_F , as a cut signature of G^* , is equal to σ_{F^*} , the F^* -connected signature of G^* (cf. Example 5.1.5). In particular, σ_F is acyclic, and any Bernardi bijection of G is a geometric bijection coming from the ABKS decomposition.*

PROOF: Without loss of generality, F is the unbounded face of G , so σ_F orients every C into clockwise. For any such directed cycle C , the dual edge of every edge $e \in C$ will be

oriented away from F , which makes C^* a directed cut oriented from away from F^* in G^* .
 \square

Next, we give alternative proofs of several properties of planar Bernardi bijections, which were proven using combinatorial arguments in [13]. Recall from Section 5.6 that every bijection $\beta : \mathcal{ST}(G) \rightarrow \text{Pic}^k(G)$ induces a group action \cdot_β of $\text{Jac}(G) \cong \text{Pic}^0(G)$ on $\mathcal{ST}(G)$ by $g \cdot_\beta T := \beta^{-1}(g + \beta(T))$. We call those group actions induced by Bernardi bijections the *Bernardi torsors*.

Theorem 6.2.4 [13, Theorem 5.1] *Let G be a planar ribbon graph. Then all Bernardi torsors are isomorphic.*

We give two proofs of Theorem 6.2.4. While the first proof is more direct and gives a more general result, the second proof provides a rather unexpected connection between Bernardi processes and the ABKS decomposition.

Proposition 6.2.5 *Let M be a regular matroid. Fix an arbitrary total ordering $e_1 < \dots < e_m$ and an arbitrary reference orientation of $E(M)$, and let σ be the acyclic circuit signature induced by such data (cf. Example 5.1.4). Define another acyclic circuit signature σ' using the same ordering of $E(M)$ and reference orientation of e_2, \dots, e_m , but with an opposite orientation of e_1 . Then $\beta_{\sigma, \tau}, \beta_{\sigma', \tau}$ induce isomorphic $\text{Jac}(M)$ -torsors for any acyclic cocircuit signature τ . In particular, when specified to cycle orientation maps, $\beta_\sigma, \beta_{\sigma'}$ induce isomorphic torsors.*

PROOF: For every basis B of M , we compare the orientation $\mathcal{O} := \hat{\beta}_{\sigma, \tau}(B)$ and $\mathcal{O}' := \hat{\beta}_{\sigma', \tau}(B)$, here $\hat{\beta}$'s are the bases to orientations maps defined in Theorem 5.4.2.

Case I: $e_1 \notin B$. e_1 is the only element whose fundamental circuit (which respect to B) involves e_1 , hence the only difference between \mathcal{O} and \mathcal{O}' is the orientation of e_1 .

Case II: $e_1 \in B$. The orientation of e_1 remains the same in $\mathcal{O}, \mathcal{O}'$ as it only depends on the fixed cocircuit signature τ . The fundamental circuit of an element $f \notin B$ involves e_1 if

and only if f is in the fundamental cocircuit $K := C^*(B, e_1)$ of e_1 , hence \mathcal{O} and \mathcal{O}' differ precisely by the orientation of these elements. Both $\hat{\beta}_{\sigma, \tau}$ and $\hat{\beta}_{\sigma', \tau}$ orient the elements of $K \setminus \{e_1\}$ consistently opposite to the reference orientation of e_1 with respect to K , so K is a positive cocircuit of either \mathcal{O} or \mathcal{O}' . And \mathcal{O} and \mathcal{O}' differ by a cocircuit reversal of K followed by a reversal of e_1 .

Summarizing the above two cases, $\beta_{\sigma, \tau}(B)$ and $\beta_{\sigma', \tau}(B)$ always differ by $[e_1] \in \frac{C_1(M)}{Z_1(M) \oplus B_1(M)} \cong \text{Jac}(M)$. By Lemma 5.6.2, they induce isomorphic torsors. \square

FIRST PROOF OF THEOREM 6.2.4: We first claim that if q, q' are two adjacent vertices in a graph G , then the q -connected signature σ_q and the q' -connected signature $\sigma_{q'}$ can be induced by data satisfying the assumption in Proposition 6.2.5. We follow the recipe described in Example 5.1.5. First we pick any spanning tree T containing an edge e incident to both q and q' , then we pick a depth-first-search ordering of edges in T starting with e from q , it is easy to see that such ordering also extends the tree ordering of edges rooted at q' . Concerning the reference orientation of edges, e will be oriented differently in the two cases, but every other edge in T will be oriented in the same way.

By Corollary 6.2.2, every Bernardi bijection is a cycle orientation map of the form β_{σ_F} , in which σ_F can be thought as a F^* -connected signature of G^* by Theorem 6.2.3. Since G^* is connected, repeatedly applying Proposition 6.2.5 yields the desired conclusion. \square

SECOND PROOF OF THEOREM 6.2.4: Fix an arbitrary vertex q of G . By Corollary 5.7.3 and Corollary 6.2.2, every Bernardi bijection can be interpreted as $\beta_{\sigma_F, \sigma_q}$ for some face F of G , where σ_F can be thought as a F^* -connected signature of G^* by Theorem 6.2.3. Now by the dual of Theorem 5.6.3, $\beta_{\sigma_F, \sigma_q}$ and $\beta_{\sigma_{F'}, \sigma_q}$ induce isomorphic group actions if and only if the two extension tilings of Z_G , induced by $\sigma_F, \sigma_{F'}$, respectively, are equivalent up to translation. But by Theorem 5.7.2, any of these extension tilings is equivalent to the ABKS decomposition of G^* lifted to its universal cover. \square

The second proposition from [13] that we will be studying is the following duality result.

Theorem 6.2.6 [13, Theorem 6.1] *Let β, β^* be some Bernardi bijections on G and G^* , respectively. Then the following diagram commutes:*

$$\begin{array}{ccc} \text{Jac}(G) \times \mathcal{ST}(G) & \xrightarrow{\cdot\beta} & \mathcal{ST}(G) \\ \downarrow \alpha_0 \times \alpha_T & & \downarrow \alpha_T \\ \text{Jac}(G^*) \times \mathcal{ST}(G^*) & \xrightarrow{\cdot\beta^*} & \mathcal{ST}(G^*) \end{array}$$

Here the map α_T from $\mathcal{ST}(G)$ to $\mathcal{ST}(G^*)$ is the map $T \mapsto E(G) \setminus T$, and the map α_0 from $\text{Jac}(G) \cong \frac{C_1(G)}{Z_1(G) \oplus B_1(G)}$ to $\text{Jac}(G^*) \cong \frac{C_1(G^*)}{Z_1(G^*) \oplus B_1(G^*)}$ is the map induced by the isomorphisms between $C_1(G)$ and $C_1(G^*)$ (by $\sum e_i \mapsto \sum e_i^*$), $Z_1(G)$ and $B_1(G^*)$, and $B_1(G)$ and $Z_1(G^*)$ [3, Proposition 8].

We prove Theorem 6.2.6 by proving the commutativity of two finer diagrams separately. In particular, our proof produces a stronger assertion than the proof in [13].

Proposition 6.2.7 *The following diagram commutes:*

$$\begin{array}{ccc} \text{Pic}^0(G) \times \text{Pic}^{g-1}(G) & \longrightarrow & \text{Pic}^{g-1}(G) \\ \downarrow \alpha_0 \times \alpha_{g-1} & & \downarrow \alpha_{g-1} \\ \text{Pic}^0(G^*) \times \text{Pic}^{g^*-1}(G^*) & \longrightarrow & \text{Pic}^{g^*-1}(G^*) \end{array}$$

Here $g^* = n - 1$ is the genus of G^* , the horizontal arrows correspond to the addition map, and the map α_{g-1} is $[D_O] \mapsto [D_{O^*}]$.

PROOF: Since the graph is connected, $\text{Pic}^0(G)$ is generated by elements of the form $[(v) - (u)]$, where u, v are adjacent vertices. Hence by linearity it suffices to prove

$$\alpha_0([(v) - (u)]) + \alpha_{g-1}([D]) = \alpha_{g-1}([(v) - (u)] + [D]), \forall [D] \in \text{Pic}^{g-1}(G)$$

for these $[(v) - (u)]$'s. Fix two such adjacent vertices u, v , say they are both incident to the edge e . For a divisor class $[D] \in \text{Pic}^{g-1}(G)$, we can interpret the addition $[(v) - (u)] + [D]$ as follows: pick an orientation \mathcal{O} such that $D \sim D_{\mathcal{O}}$ and that e is oriented from v to u (the latter is always possible by reversing a directed cycle/cut e is in), reverse e in \mathcal{O} to obtain \mathcal{O}' , then $[(v) - (u)] + [D_{\mathcal{O}}] = [D_{\mathcal{O}'}]$. Using the convention in Figure 6.3, the dual element of $[(v) - (u)]$ is $[(F'^*) - (F^*)] \in \text{Pic}^0(G^*)$, and the dual element of $[D] = [D_{\mathcal{O}}]$ is $[D_{\mathcal{O}^*}] \in \text{Pic}^{g^*-1}(G^*)$, note that e^* is oriented from F'^* to F^* in \mathcal{O}^* . Denote by $\mathcal{O}^{*'}$ the orientation of G^* obtained from reversing e^* in \mathcal{O}^* , then $[(F'^*) - (F^*)] + [D_{\mathcal{O}^*}] = [D_{\mathcal{O}^{*'}}]$. But it is easy to see that $\mathcal{O}^{*'}$ is the dual orientation of \mathcal{O}' , thus $[D_{\mathcal{O}^{*'}}]$ is the dual element of $[D_{\mathcal{O}'}]$. Summarizing we have the desired identity. \square

Proposition 6.2.8 *Let $\beta = \beta_F, \beta^* = \beta_{v^*}$ be Bernardi bijections of G, G^* , respectively. Then the following diagram commutes:*

$$\begin{array}{ccc} \mathcal{ST}(G) & \xrightarrow{\beta - [(v)]} & \text{Pic}^{g-1}(G) \\ \downarrow \alpha_T & & \downarrow \alpha_{g-1} \\ \mathcal{ST}(G^*) & \xrightarrow{\beta^* - [(F^*)]} & \text{Pic}^{g^*-1}(G^*) \end{array}$$

PROOF: Let T be a spanning tree of G . The divisor class $\beta(T) - [(v)]$ contains an orientable divisor $D_{\mathcal{O}}$ in which $\mathcal{O} = \hat{\beta}_{\sigma_F, \sigma_v}(T)$. Dually, the divisor class $\beta^*(E(G) \setminus T) - [(F^*)]$ contains an orientable divisor $D_{\mathcal{O}'}$ in which $\mathcal{O}' = \hat{\beta}_{\sigma_{v^*}, \sigma_{F^*}}(E(G) \setminus T)$. By symmetry, we can see that \mathcal{O}' is the dual orientation of \mathcal{O} , thus $[D_{\mathcal{O}'}] = \alpha_{g-1}([D_{\mathcal{O}}])$. \square

PROOF OF THEOREM 6.2.6: The diagram there factors through the diagram (and its reversal) in Proposition 6.2.8 and the diagram in Proposition 6.2.7 as follows.

$$\begin{array}{ccccccc} \text{Jac}(G) \times \mathcal{ST}(G) & \xrightarrow{\text{id} \times (\beta - [(v)])} & \text{Pic}^0(G) \times \text{Pic}^{g-1}(G) & \longrightarrow & \text{Pic}^{g-1}(G) & \xrightarrow{\beta^{-1}(\cdot + [(v)])} & \mathcal{ST}(G) \\ \downarrow \alpha_0 \times \alpha_T & & \downarrow \alpha_0 \times \alpha_{g-1} & & \downarrow \alpha_{g-1} & & \downarrow \alpha_T \\ \text{Jac}(G^*) \times \mathcal{ST}(G^*) & \xrightarrow{\text{id} \times (\beta^* - [(F^*)])} & \text{Pic}^0(G^*) \times \text{Pic}^{g^*-1}(G^*) & \longrightarrow & \text{Pic}^{g^*-1}(G^*) & \xrightarrow{\beta^{*-1}(\cdot + [(F^*)])} & \mathcal{ST}(G^*) \end{array}$$

\square

Finally, by [13, Theorem 7.1], Theorem 6.2.4 and 6.2.6 concerning Bernardi torsors of planar ribbon graphs can be translated to their counterparts concerning *rotor-routing torsors* [30, 31].

6.2.2 Non-planar Case

We prove a partial converse of Baker’s conjecture (which is a theorem by Section 6.2.1). That is, the Bernardi bijections of a non-planar ribbon graph are “almost never” geometric bijections from the ABKS decomposition. The intuition is that cycles of a non-planar ribbon graph need not partition the surface into “inside” and “outside”, hence do not have a consistent orientation as in the planar case.

First we characterize non-planar ribbon graphs by forbidden subdivisions. We say a ribbon graph \hat{H} is a *subdivision* of the ribbon graph H if one can obtain \hat{H} from H by inserting degree 2 vertices (equipped with the unique cyclic ordering on two edges) inside the edges of H while keeping the cyclic orderings of edges around the original vertices the same. Also, we say a ribbon graph H is a *subgraph* of a ribbon graph G if graph-theoretically H is a subgraph of G and the cyclic ordering of edges around each vertex of H is inherited from the cyclic ordering of G . Finally, we say a ribbon graph G contains a ribbon graph H as a subdivision if \hat{H} is a subgraph of G for some subdivision \hat{H} of H .

Definition 6.2.9 *The first basic non-planar ribbon graph (BNG I) is a ribbon graph on two vertices v_1, v_2 and three edges e_1, e_2, e_3 so that the cyclic ordering of the edges around each vertex is e_1, e_2, e_3 . The second basic non-planar ribbon graph (BNG II) is a ribbon graph on three vertices u_1, u_2, u_3 , with two edges e'_1, e'_2 between u_1, u_2 and two edges e'_3, e'_4 between u_1, u_3 , where the cyclic ordering of the edges around u_1 is e'_1, e'_3, e'_2, e'_4 .*

The next proposition is known within the communities working in structural or topological graph theory. But we could not find an explicit reference in the literature so we include a brief proof here.

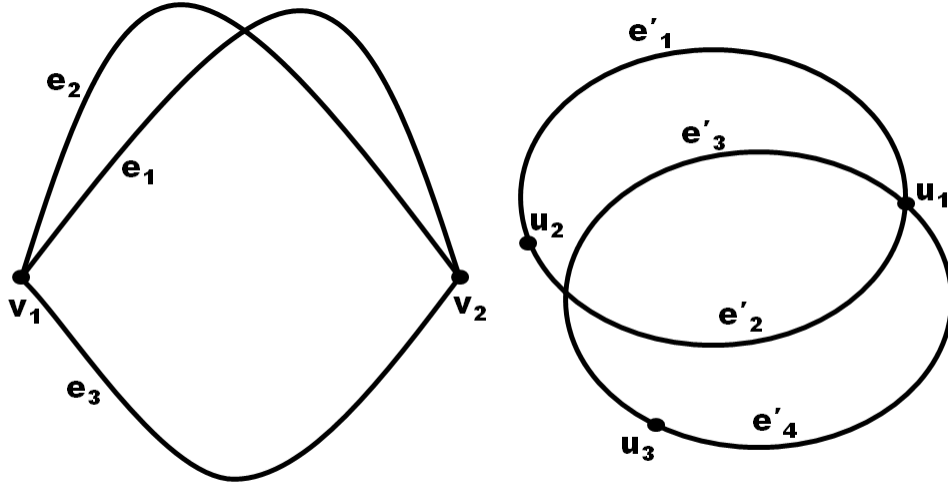


Figure 6.4: BNG I (left) and BNG II (right).

Lemma 6.2.10 *Every non-planar ribbon graph contains at least one of the BNGs as a subdivision.*

PROOF: Given a ribbon graph G , we will demonstrate a way to either find a planar embedding of G , or find a BNG as a subdivision in G . First we assume that G is 2-connected, and we apply induction via an ear decomposition. The base case is a cycle and is trivial. Suppose a subgraph $G' \subset G$ is embedded in the plane, and let $P = u_0 - e_1 - \dots - e_k - u_k$ be an ear to be added. Say in $G' \cup P$ the cyclic ordering of edges around u_0 includes e, e_1, f consecutively. Then some cycle C containing e, f bounds a face F of G' that e_1 is to be embedded in; similarly there is a face F' of G' that e_k is to be embedded in. If $F = F'$ then we can embed P in F ; otherwise if we let $Q \subset G'$ be a shortest path (possibly trivial) going from u_k to any vertex $v \neq u_0$ on C (the existence of such path is guaranteed by the 2-connectivity of G'), then $P \cup Q \cup C$ will be a subdivision of BNG I.

For the general case, we induct on the number of blocks. Let v be a cut vertex and let G_1, \dots, G_k be subgraphs corresponding to the components of $T - v$, where T is the block decomposition tree of G . By induction, each G_i can be embedded in the plane, and we may further assume that v is on the boundary of the unbounded face of each embedding if

needed. If there exist some interlacing edges e, e', f, f' in the cyclic ordering around v with $e, f \in G_i$ and $e', f' \in G_j$, then by letting $C \subset G_i, C' \subset G_j$ to be cycles containing e, f and e', f' , respectively, we have $C \cup C'$ as a subdivision of BNG II in G . Otherwise, it can be seen that for all subgraphs G_i 's except possibly one (say G_1), all edges in G_i incident to v are in some interval I_i of the cyclic ordering around v , so we can embed G_1 in the plane and then embed other subgraphs one by one according to the cyclic ordering of I_i 's. \square

For the proof we adopt the following conventions: we say an edge $f \neq e_1, e_2$ in a ribbon graph is *in between* e_1 and e_2 at v if v is a common endpoint of the three edges, and f goes before e_2 in the cyclic ordering of edges around v when listed starting with e_1 ; given a simple path P and vertices $a, b \in V(P)$, we denote by aPb the subpath of P between a and b (inclusive); and given a spanning tree T and a subset V' of vertices such that $T[V']$ is connected, we say a vertex v not in V' is *under* a vertex $v' \in V'$ if the closest vertex from V' to v in T is v' .

Theorem 6.2.11 *Let G be a non-planar ribbon graph. If G is simple, then there exists some Bernardi bijection of G that is not a cycle orientation map. Otherwise there exists some subdivision G' of G such that some Bernardi bijection of G' is not a cycle orientation map, and G' can be chosen to have at most one more vertex than G .*

PROOF: Let G be a non-planar ribbon graph containing a subdivision $P_1 \cup P_2 \cup P_3$ of BNG I, with P_1, P_2, P_3 being internally disjoint paths sharing endpoints v_1, v_2 ; we assume P_1 is not an edge in the simple graph case by re-indexing. Without loss of generality, we may assume there are no edges between (the last edge of) P_1 and (the last edge of) P_2 whose endpoints are v_2 and some internal vertex of P_1 or P_2 : if there is such an edge $f = v_2t$ with $t \in V(P_i), i = 1$ or 2 , then we can replace P_i by $v_1P_it \cup \{f\}$, the process will eventually stop because the number of edges between P_1 and P_2 decreases in every step. Note that in the case of simple graphs, the process will keep at least one internal vertex of P_1 . Similarly we may assume there are no edges between v_1 and v_2 that are between P_1, P_2 at the two

ends, or otherwise we may replace P_2 by such edge. By inserting a new vertex on P_1 near v_2 in the non-simple case if necessary, we may assume $P_1 = v_1 - \dots - e_{11} - u - e_{12} - v_2$ is of length at least 2, and there are no edges between u and v_2 other than e_{12} .

Denote by e_3 the edge on P_3 that is incident to v_2 , we extend the acyclic subgraph $(P_1 - e_{11}) \cup P_2 \cup (P_3 - e_3)$ to a spanning tree T_1 of G with the maximum number of vertices under v_2 with respect to $V[P_1 \cup P_2 \cup P_3]$. Note that our assumption means the other endpoint of any non-tree edge incident to v_2 in T_1 is either from $V[P_1 \cup P_2 \cup P_3]$ or is a vertex under v_2 . Set $T_2 = T_1 - e_{12} + e_{11}$. It is easy to see that the set of non-tree edges incident to v_2 in T_2 is exactly the set of non-tree edges incident to v_2 in T_1 plus e_{12} , and those common non-tree edges have the same fundamental cycles in T_1 and T_2 . Consider the Bernardi tours of $T_1, T_2 \subset G$ starting from (u, e_{11}) . A routine simulation shows that each non-tree edge that contributes a chip to v_2 in the first tour will also contribute a chip to v_2 in the second, while e_{12} and e_3 will each contribute a chip to v_2 in the second tour but not in the first, so v_2 received at least two more chips in the second tour than in the first. But in any cycle orientation map, v_2 can receive at most one more chip with respect to T_2 than T_1 , so the Bernardi bijection is not a cycle orientation map.

The case when G only has subdivisions of BNG II is similar. Let $C_1 = u_1 - e_1 - u - f - \dots - e_3 - u_1$ (f could be equal to e_3) and $C_2 = u_1 - e_2 - \dots - w - e_4 - u_1$ be two cycles of G whose union is a subdivision of BNG II, i.e., the two cycles are disjoint except at u_1 , and the cyclic ordering of edges around u_1 includes e_1, e_2, e_3, e_4 in order. With a greedy procedure similar to the one in the case of BNG I, we may assume there are no edges between e_1 and e_2 whose endpoints are u_1 and some other vertex of C_1 or C_2 , nor edges between e_2 and e_3 whose endpoints are u_1 and some other vertex of C_1 . By inserting a new vertex in e_1 near u_1 in the non-simple case, we may further assume there are no edges between u_1 and u other than e_1 . Extend $(C_1 - f) \cup (C_2 - e_4)$ to a spanning tree T'_1 of G with the maximum number of vertices under u_1 with respect to $V[C_1 \cup C_2]$, and set $T'_2 = T'_1 - e_1 + f$. Consider the two Bernardi tours of T'_1, T'_2 starting from (u, f) .

u_1 received at least two more chips (from e_1 and e_4) in the second tour than in the first, a similar reasoning as above shows the Bernardi bijection cannot be a cycle orientation map. \square

Example 6.2.12 *It can be checked that all Bernardi bijections on BNG I are coming from its ABKS decomposition, so working with subdivisions is indeed necessary in Theorem 6.2.11.*

6.3 Bernardi Bijections as Burning Algorithms

All Bernardi bijections in this section are of the second type, i.e., they map spanning trees to v -reduced divisors. We interpret the Bernardi bijections as a variant of Dhar's burning algorithm (Proposition 2.2.8). Such idea was known in Bernardi's original paper, but our proof seems new. In particular, it uses idea closely related to the inverse of an edge ordering map discussed in Section 5.8.3 as well as the proof of Proposition 6.2.5. In contrast to the last section, the result here does not require the ribbon graph to be planar.

6.3.1 Cori–Le Borgne Processes

Cori and Le Borgne [34] gave a bijective proof to a theorem of Merino [86], which essentially says that the statistics of q -reduced divisors can be used to characterize $T_G(1, y)$. Formally, Merino's theorem can be stated as

$$T_G(1, y) = y^g \sum_D y^{D(q)},$$

where the sum runs over all q -reduced divisor D of $\text{Div}^0(G)$. To achieve this, they modified Dhar's algorithm into a bijection from q -reduced divisors to spanning trees by introducing a tiebreaking rule to the algorithm, namely a total ordering $<$ of all edges. Instead of letting the fire spread in an arbitrary order, in each step the largest unburnt edge that has exactly one burnt end is set on fire, and if setting an edge f on fire causes the unburnt end of f to

be burnt, include f into the output edge set. Cori and Le Borgne proved that the algorithm always produces a spanning tree, and it is a bijection. Moreover, the edges not being burnt in the end of the algorithm are externally active with respect to the output tree T and $<$, that is, $\epsilon(T) = g + D(q)$.

Input: A q -reduced divisor $D \in \text{Div}(G)$, and a total ordering $<$ of E .
Output: A spanning tree T of G .

Set $X := \{q\}$, $R := \emptyset$, $T := \emptyset$. (Burnt vertices, burnt edges and marked edges, respectively)

while $X \neq V$ **do**

$f := \max(G[X, \overline{X}] \cap \overline{R})$

$v :=$ the unburnt vertex incident to f

if $D(v) = |R \cap \delta(v)|$ **then**

$X := X \cup \{v\}$

$T := T \cup \{f\}$

end

$R := R \cup \{f\}$

end

Output T .

Algorithm 5: Cori–Le Borgne Algorithm

6.3.2 Statement and Proof

Theorem 6.3.1 *Fix a ribbon structure and a starting pair (v, e) of G . Let T be a spanning tree of a graph G . Perform Bernardi process on T to obtain a v -reduced divisor D as well as a total ordering $<_T$ of edges. Then performing Cori–Le Borgne process on D using $<_T$ produces back T .*

PROOF: Apply induction on m . The case with no edges is vacuously true. So we assume $m > 0$.

Case I: e is not in T . Let f be the edge following e around v . Consider a new Bernardi process on T as a spanning tree of $G - e$ with starting edge (v, f) . Say the divisor and total ordering the new process produces are D' , $<'_T$, respectively. Since $e \notin T$, the new Bernardi tour is the restriction of the original Bernardi tour on $G - e$. From this we can see that $<'_T$ is

the same as $<_T$ restricted to $E - e$; also we have $D' = D$ because e is externally Bernardi active hence does not contribute any chips to D . Performing Cori–Le Borgne process on D as a divisor of $G - e$ using $<'_T$ produces back T by induction, while performing the process on D as a divisor of G using $<_T$ will follow the same initial steps, and before the algorithm considers e (the only difference between the two instances), the algorithm would terminate because a spanning tree (namely T) is already found.

Case II: e is in T . Let T_1, T_2 be the two components of $T - e$, without loss of generality $v \in T_2$ and say the other vertex of e is $u \in T_1$. Let $V_1 := V(T_1), V_2 := V(T_2), G_1 := G[V_1], G_2 := G[V_2]$. Let $K := G[V_1, V_2]$ be the fundamental cut of e and write $K' := K - e$. Let f be the edge following e around v and let g be the edge following e around u . The Bernardi process on T will first visit e , then visit all edges in G_1 (twice) as well as all edges in K' (once each), next it will visit e the second time, and finally visit edges in G_2 (twice) as well as edges in K' for the second time each. In particular, every edge in G_2 is larger than any edge in $E(G_1) \cup K$ with respect to $<_T$, and all chips contributed by K' are in G_1 . Notice that the smallest edge e is contained in the fundamental cycle of every edge in K' , so every edge in K' is not Bernardi externally active and will contribute a chip. Therefore $D(x) \geq \text{outdeg}_{V_1}(x)$ for every $x \in V_1 \setminus \{u\}$ and $D(u) \geq \text{outdeg}_{V_1}(u) - 1$ (hence by the v -reducedness of D , $D(u) = \text{outdeg}_{V_1}(u) - 1$).

Consider two smaller Bernardi processes. One on T_1 as a spanning tree of G_1 , starting with (u, g) ; and another on T_2 as a spanning tree of G_2 , starting with (v, f) . Say the divisors and total orderings produced by these two new processes are $D_1, <^1_T$ and $D_2, <^2_T$, respectively. The Bernardi tours of the two small processes are the restriction of the original Bernardi tour on G_1 and G_2 , respectively. Therefore $<^1_T$ (resp. $<^2_T$) is $<_T$ restricted to $E(G_1)$ (resp. $E(G_2)$). We also have that $D|_{V_2}$ equals D_2 , while $D|_{V_1}$ equals D_1 plus a chip from each edge in K' .

Now perform Cori–Le Borgne process on D (as a divisor of G) with $<_T$. From the observations regarding $D(x)$ and $\text{outdeg}_{V_1}(x)$ for $x \in V_1$, we know that as long as e is

not burnt, none of the vertices in G_1 would be burnt. Therefore the burning process will proceed in three phrases:

1. At the beginning, the process only needs to burn along edges of G_2 in order to burn out every vertex in V_2 , which produces T_2 , both assertions follow from induction hypothesis on G_2 .
2. Now the set of edges with exactly one end being burnt is precisely K . Since e is the smallest edge in K while $D(x) \geq \text{outdeg}_{V_1}(x), \forall x \in V_1 \setminus \{u\}$, all edges in K will be burnt and the first burnt vertex in V_1 is u . This phase essentially exhausts all firefighters coming from K' during the Bernardi process and reduces the configuration of available firefighters in G_1 from $D|_{V_1}$ to D_1 .
3. Finally by induction hypothesis on G_1 , the process will burn through all vertices in V_1 along edges in G_1 to produce T_1 .

To summarize, the Cori–Le Borgne process will give back $T_1 \cup \{e\} \cup T_2 = T$ as claimed.

□

In the BIRS–CMO Workshop on sandpile groups, Hopkins asked whether one can classify all tiebreak rules for Dhar’s burning algorithm that yield bijections between spanning trees and reduced divisors, as well as the corresponding statistics of spanning trees [68, Question 1.4]. Our result provides another set of examples to Hopkins’ question.

Corollary 6.3.2 *Fix a ribbon structure and a starting pair (v, e) of G . Then using $<_T$ as the tiebreak rule for each (degree 0) v -reduced divisor D in the burning algorithm yields a bijection, here $T := T(D)$ is the unique spanning tree that corresponds to D in the Bernardi bijection. Furthermore, such bijection relates the value of D at q and the external Bernardi activity e_B of T : we have $e_B(T) = g + D(v)$.*

The fact that both classical activity and Bernardi activity are special cases of Δ -activity introduced by Courtiel [35] suggests there might be one variant of burning algorithms for every Δ -activity, which would provide a family of examples to Hopkins’ question.

CHAPTER 7

EXTENSIONS

7.1 Geometric Bijections for General Oriented Matroids

This section is based on joint work with Spencer Backman and Francisco Santos.

We will prove a generalization of Theorem 5.4.2 to include all oriented matroids. We first start with the necessary background to state and prove the theorem.

Definition 7.1.1 *Let M be an oriented matroid on ground set E . An oriented matroid M' is a single-element extension of M if the ground set of M' is $E \sqcup \{f\}$ for some new element f and $M = M' \setminus f$. Dually, \widetilde{M} is a single-element lifting of M if the ground set of \widetilde{M} is $E \sqcup \{g\}$ and $M = \widetilde{M}/g$.*

Let M' be a single-element extension of M . By the definition of oriented matroid deletion, for every signed cocircuit D of M , there exists a unique signed cocircuit D' of M' such that $D'|_E = D$. Therefore we can define a *cocircuit signature of signs* $\sigma^* : \mathcal{C}^*(M) \rightarrow \{+, 0, -\}$ associated to the extension by setting $D \mapsto D'(f)$. Dually, every single-element lifting is associated with a *circuit signature of signs* $\sigma : \mathcal{C}(M) \rightarrow \{+, 0, -\}$.

We have a converse uniqueness statement.

Theorem 7.1.2 [23, Proposition 7.1.4] *If M', M'' are single-element extensions of M on the same ground set that induce the same cocircuit signature of signs, then $M' = M''$. A dual statement for single-element liftings.*

In general, not every circuit (resp. cocircuit) signature of signs is coming from an actual single-element lifting (resp. single-element extension). For instance, we must have $\sigma(-C) = -\sigma(C)$ for every signed circuit C . A list of necessary and sufficient conditions is given in [23, Theorem 7.1.8].

The definition of circuit (resp. cocircuit) signatures in Chapter 5 is compatible with the definition here. We say a signature of signs, as well as the single-element lifting or extension associating to the signature, is *generic* if the range is $\{+, -\}$. Every generic circuit signature of signs σ induces a signature $\underline{\sigma} : \mathcal{C}(\underline{M}) \rightarrow \mathcal{C}(M)$ in the way that $\sigma(\underline{\sigma}(C)) = +$ for every $C \in \mathcal{C}(\underline{M})$. Conversely, every signature $\underline{\sigma} : \mathcal{C}(\underline{M}) \rightarrow \mathcal{C}(M)$ induces a generic circuit signature of signs $\sigma : \mathcal{C}(M) \rightarrow \{+, -\}$ by setting $\sigma(C) = +$ if and only if $C = \underline{\sigma}(C)$.

We will show that using the dictionary above, every acyclic circuit (resp. cocircuit) signature is equivalent to a circuit (resp. cocircuit) signature of signs associated to certain single-element lifting (resp. extension), and such lifting (resp. extension) can be described explicitly.

Proposition 7.1.3 *Let M be an oriented matroid realized by the $r \times m$ matrix A , which we assume is of full row rank. Let $\mathbf{w} \in \mathbb{R}^E$ be a generic vector in the sense of Lemma 5.3.2. We have that*

1. *The oriented matroid M' realized by*

$$A' = \begin{pmatrix} A & A\mathbf{w} \end{pmatrix}$$

is a single-element extension of M . Let $\sigma^ : \mathcal{C}^*(M) \rightarrow \{+, 0, -\}$ be the associated cocircuit signature of signs. Then σ^* is generic, and the map $\underline{\sigma}^* : \mathcal{C}^*(\underline{M}) \rightarrow \mathcal{C}^*(M)$ induced by σ^* coincides with the acyclic cocircuit signature induced by \mathbf{w} in the way described in Lemma 5.3.2.*

2. *The oriented matroid \widetilde{M} realized by*

$$\widetilde{A} = \begin{pmatrix} A & \mathbf{0} \\ \mathbf{w}^T & -1 \end{pmatrix}$$

is a single-element lifting of M . Let $\sigma : \mathcal{C}(M) \rightarrow \{+, 0, -\}$ be the associated circuit

signature of signs. Then σ is generic, and the map $\underline{\sigma} : \mathcal{C}(\underline{M}) \rightarrow \mathcal{C}(M)$ induced by σ coincides with the acyclic circuit signature induced by \mathbf{w} .

PROOF: Let D be a signed cocircuit of M . By definition, D is the sign pattern of $\mathbf{u}^T A \neq \mathbf{0}$ for some $\mathbf{u} \in \mathbb{R}^r$. We claim that the sign pattern D' of $\mathbf{u}^T A' = \begin{pmatrix} \mathbf{u}^T A & \mathbf{u}^T A \mathbf{w} \end{pmatrix}$ is a cocircuit of M' : suppose the support of $\mathbf{u}^T A'$ is properly contained in the support of D' , then $\mathbf{u}^T A \neq \mathbf{0}$ and its support is properly contained in the support of D , contradicting the assumption that D is a cocircuit of M . Conversely, if D' is a signed cocircuit of M' , realized as the sign pattern of $\mathbf{u}^T A'$, then $\mathbf{u}^T A \neq \mathbf{0}$ is a signed covector of M . Hence M' is a single-element extension of M . Furthermore, by the generic assumption on \mathbf{w} , $\mathbf{u}^T A \mathbf{w} \neq 0$ whenever $\mathbf{u}^T A$ realizes some signed cocircuit, so σ^* is generic. Finally, we have $D = \underline{\sigma}^*(\underline{D})$ if and only if $\sigma^*(D) = +$, if and only if the last coordinate of D' is positive, if and only if $\mathbf{u}^T A \mathbf{w} > 0$, which is the condition given in Lemma 5.3.2.

The dual statement can be proven in a similar manner. If C is a signed circuit of M , then it is the sign pattern of some $\mathbf{v} \in \ker A$, and $\begin{pmatrix} \mathbf{v} \\ \mathbf{w}^T \mathbf{v} \end{pmatrix}$ realizes a signed circuit \tilde{C} of \tilde{M} . And $C = \underline{\sigma}(\underline{C})$ if and only if $\sigma(C) = +$, if and only if the last coordinate of \tilde{C} is positive, if and only if $\mathbf{w}^T \mathbf{v} > 0$. \square

Now we state the generalization of Definition 5.4.1 and Theorem 5.4.2.

Definition 7.1.4 *Let M be an oriented matroid on ground set E . Let M' , \tilde{M} be a generic single-element extension and a generic single-element lifting of M on ground sets $E \sqcup \{f\}$ and $E \sqcup \{g\}$, respectively. Let σ^* (resp. σ) be the cocircuit (resp. circuit) signature of signs associated to M' (resp. \tilde{M}). Then an orientation \mathcal{O} of M is (σ, σ^*) -compatible if every signed circuit (resp. cocircuit) compatible with \mathcal{O} is oriented according to $\underline{\sigma}$ (resp. $\underline{\sigma}^*$). Then set of (σ, σ^*) -compatible orientations of M is denoted as $\chi(M; \sigma, \sigma^*)$.*

For an orientation \mathcal{O} of M , \mathcal{O}'_- is the orientation of M' such that $\mathcal{O}'_-|_E = \mathcal{O}$ and $\mathcal{O}'_-(f) = -$; dually, $\tilde{\mathcal{O}}_-$ is the orientation of \tilde{M} such that $\tilde{\mathcal{O}}_-|_E = \mathcal{O}$ and $\tilde{\mathcal{O}}_-(g) = -$.

For the rest of this section, $M, M', \widetilde{M}, \sigma, \sigma^*$ will always be the same as in Definition 7.1.4, unless otherwise specified.

Theorem 7.1.5 *Given a basis $B \in \mathcal{B}(M)$, let $\mathcal{O}(B)$ be the orientation of M in which we orient each $e \notin B$ according to its orientation in $\underline{\sigma}(C(B, e))$ and each $e \in B$ according to its orientation in $\underline{\sigma}^*(C^*(B, e))$. Then the map $B \mapsto \mathcal{O}(B)$ gives a bijection $\hat{\beta}_{\sigma, \sigma^*} : \mathcal{B}(M) \rightarrow \chi(M; \sigma, \sigma^*)$.*

In view of Proposition 7.1.3, Theorem 7.1.5 is indeed a generalization of Theorem 5.4.2.

Theorem 7.1.5 will be deduced from the following theorem, which was mainly formulated and proved by Santos.. Recall from Section 4.1 that a circuit (resp. cocircuit) $C \subset E$ is positive with respect to an orientation \mathcal{O} if one of the signed versions of C is compatible with \mathcal{O} .

Theorem 7.1.6 *For every $\mathcal{O} \in \chi(M; \sigma, \sigma^*)$, there exists a unique basis $B \in \mathcal{B}(M)$ such that $B \cup \{f\}$ is a positive circuit of \mathcal{O}'_- and $(E \setminus B) \cup \{g\}$ is a positive cocircuit of $\widetilde{\mathcal{O}}_-$.*

We start with a few lemmas.

Lemma 7.1.7 [97, Lemma 1.10] *There exists an oriented matroid \widetilde{M}' on ground set $E \sqcup \{f, g\}$ such that $M' = \widetilde{M}'/g$ and $\widetilde{M} = \widetilde{M}' \setminus f$.*

Lemma 7.1.8 *Let M be a matroid on ground set E and let M' be a generic extension of M on ground set $E \sqcup \{f\}$. Then the set of circuits of M' containing f is $\{B \cup \{f\} : B \in \mathcal{B}(M)\}$. Dually, if \widetilde{M} is a generic lifting of M on ground set $E \sqcup \{g\}$, then the set of cocircuits of \widetilde{M} containing g is $\{(E \setminus B) \cup \{g\} : B \in \mathcal{B}(M)\}$.*

PROOF: Let $B \in \mathcal{B}(M)$. We first claim that B is also a basis of M' . Since every circuit of M' not containing f is a circuit of M , B is independent in M' ; since every circuit of M is a circuit of M' , $B \cup \{e\}$ is dependent in M' for any $e \in E \setminus B$. So if B is not a basis of M' , it must be the case that $X := B \cup \{f\}$ is a basis of M' . In such case, $B = X \setminus \{f\}$

avoids the fundamental cocircuit D' of f with respect to X in M' . Since M' is generic, f is not an isthmus and $D' \setminus \{f\}$ contains a cocircuit D'' of M , now B avoids the cocircuit D'' in M , contradicting the basic property of bases.

Next we claim that the fundamental circuit C' of f with respect to B is the whole of X . Suppose not, pick an arbitrary $e \in X \setminus C'$ and let D be the fundamental cocircuit of e with respect to B in M . On one hand, $D' := D \cup \{f\}$ is a cocircuit of M' as the extension is generic, so D' must be the fundamental cocircuit of e with respect to B in M' . On the other hand, since $e \notin C' = C(B, f)$, f cannot be in $D' = C^*(B, e)$, a contradiction. This shows $\{B \cup \{f\} : B \in \mathcal{B}(M)\} \subset \mathcal{C}(M')$.

Conversely, let $C' \in \mathcal{C}(M')$ be a circuit containing f . Then $Y := C' \setminus \{f\}$ is independent in M' thus in M . If Y is not a basis of M , then it is properly contained in some $B \in \mathcal{B}(M)$, but by the above containment, $B \cup \{f\}$ is a circuit of M' properly containing C' , a contradiction. The dual statement can be proven similarly. \square

Lemma 7.1.9 *An orientation \mathcal{O} of M is σ^* -compatible if and only if \mathcal{O}'_- is totally cyclic. Dually, \mathcal{O} is σ -compatible if and only if $\tilde{\mathcal{O}}_-$ is acyclic.*

PROOF: Suppose \mathcal{O}'_- is compatible with some signed cocircuit D' . By [23, Proposition 7.1.4 (ii)], $D := D'|_E$ is either (i) a signed cocircuit of M , in which $f \in \underline{D}$, or (ii) equal to the conformal composition $D_1 \circ D_2$ of signed cocircuits of M , in which $\sigma^*(D_1) = -\sigma^*(D_2) \neq 0$. For case (i), D is a signed cocircuit compatible with \mathcal{O} , but it is not compatible with $\underline{\sigma}^*$ as $D'(f) = \mathcal{O}'_-(f) = -$; for case (ii), both D_1, D_2 are compatible with \mathcal{O} , but at least (exactly) one of them is not compatible with $\underline{\sigma}^*$ as $\sigma^*(D_1) = -\sigma^*(D_2)$.

Conversely, if D is a signed cocircuit compatible with \mathcal{O} but not $\underline{\sigma}^*$, then $(D -)$ is a signed cocircuit of M' that is compatible with \mathcal{O}'_- , hence \mathcal{O}'_- is not totally cyclic. The dual statement can be proven similarly. \square

Next we clarify the relation between several classes of orientations which have appeared so far.

Proposition 7.1.10 *Let B be a basis of M and let $\mathcal{O} = \hat{\beta}_{\sigma, \sigma^*}(B)$. Then $B \cup \{f\}$ is a positive circuit of \mathcal{O}'_- and $(E \setminus B) \cup \{g\}$ is a positive cocircuit of $\tilde{\mathcal{O}}_-$. This gives an alternative description of the map $\hat{\beta}_{\sigma, \sigma^*}$.*

PROOF: By Lemma 7.1.8, $X := B \cup \{f\}$ is a circuit of \underline{M}' . Denote by C the signed circuit of M' whose support is X and satisfies $C(f) = -$. For every $e \in B$, let D_e be the fundamental cocircuit of e with respect to B in M , oriented according to $\underline{\sigma}^*$. By the definition of $\underline{\sigma}^*$, the signed subset $D'_e := (D_e +)$ is a signed cocircuit of M' , and $X \cap \underline{D}'_e = \{e, f\}$. By orthogonality of signed circuits and cocircuits as well as the fact that $D'_e(f) = -C(f)$, we must have $\mathcal{O}(e) = D_e(e) = D'_e(e) = C(e)$. Therefore X is oriented as C in \mathcal{O}'_- and thus a positive circuit. The second statement is the dual of the first one. \square

Proposition 7.1.11 *Let \mathcal{O} be an orientation of M . If there exists a basis $B \in \mathcal{B}(M)$ such that $B \cup \{f\}$ is a positive circuit of \mathcal{O}'_- and $(E \setminus B) \cup \{g\}$ is a positive cocircuit of $\tilde{\mathcal{O}}_-$, then $\mathcal{O} \in \mathcal{X}(M; \sigma, \sigma^*)$.*

PROOF: By Lemma 7.1.9, it suffices to show that \mathcal{O}'_- is totally cyclic and $\tilde{\mathcal{O}}_-$ is acyclic. Suppose D is a signed cocircuit compatible with \mathcal{O}'_- . Since B is also a basis of M' (cf. the proof of Lemma 7.1.8), $X := \underline{D} \cap B$ is non-empty, but then X will be simultaneously in the circuit part and cocircuit part of \mathcal{O}'_- , contradicting Proposition 4.1.1. The dual statement can be proven similarly. \square

The proof of Theorem 7.1.6 relies on the theory of *oriented matroid programming*, we assume the basic terminology and results from [23, Chapter 10].

PROOF OF THEOREM 7.1.6: “Uniqueness”. Suppose both B_1 and B_2 are bases satisfying the condition. Let C_1, C_2 be the signed circuits of M' obtained from restricting \mathcal{O}'_- to $B_1 \cup \{f\}$ and $B_2 \cup \{f\}$, respectively; let D_1, D_2 be the signed cocircuits of \tilde{M} obtained from restricting $\tilde{\mathcal{O}}_-$ to $(E \setminus B_1) \cup \{g\}$ and $(E \setminus B_2) \cup \{g\}$, respectively. Let \tilde{M}' be the

oriented matroid containing both M' and \widetilde{M} as guaranteed by Lemma 7.1.7 and consider the lift \widetilde{C}_1 of C_1 in \widetilde{M}' .

Case I: $\widetilde{C}_1(g) = +$. Let D'_1, D'_2 be the extensions of D_1, D_2 in \widetilde{M}' . We must have $D'_1(f) = D'_2(f) = -$ by orthogonality, which in turn forces the lift \widetilde{C}_2 of C_2 to take value $+$ at g . Apply the circuit elimination axiom to \widetilde{C}_1 and $-\widetilde{C}_2$ and eliminate f . Denote by C the resulting signed circuit. We have $\underline{C} \cap \underline{D}'_1 \subset (B_2 \setminus B_1) \cup \{g\}$, but C is conformal with $-D'_1$ over $B_2 \setminus B_1$ as $D'_1|_{B_2 \setminus B_1} = \mathcal{O}|_{B_2 \setminus B_1} = C_2|_{B_2 \setminus B_1}$, so $C(g) = D'_1(g) = -$ by orthogonality. However, the same orthogonality argument applied to C and D'_2 implies that $C(g) = -D'_2(g) = +$, a contradiction.

Case II: $\widetilde{C}_1(g) = -$. The analysis is similar to Case I.

Case III: $\widetilde{C}_1(g) = 0$. This case is impossible as well, as \widetilde{C}_1 cannot be orthogonal to D'_1, D'_2 in the first place.

“Existence”. Let $\mathcal{O} \in \chi(M; \sigma, \sigma^*)$. By reorienting M if necessary, we may assume $\mathcal{O} \equiv +$. For the sake of matching convention in the literature, we also reorient f, g in \widetilde{M}' , so the all positive orientation \mathcal{O}'_+ of M' is totally cyclic and the all positive orientation $\widetilde{\mathcal{O}}_+$ is acyclic by Lemma 7.1.9. Now we consider the *oriented matroid program* $\mathcal{P} := (\widetilde{M}', g, f)$.

\mathcal{P} is both *feasible* and *bounded* from our assumption on $\widetilde{\mathcal{O}}_+$ and \mathcal{O}'_+ : $\widetilde{\mathcal{O}}_+$ itself is a positive covector of \widetilde{M} , which corresponds to a (full-dimensional) feasible region; any positive circuit of M' whose support is of the form $B \cup \{f\}$, $B \in \mathcal{B}(M)$ provides a bounded cone B containing the feasible region. By the main theorem of oriented matroid programming [23, Theorem 10.1.13], \mathcal{P} has an optimal solution Y , which is a covector of \widetilde{M}' .

By definition, Y is *feasible* and *optimal*, i.e., $Y(g) = +$, $Y|_E \geq 0$, and $Y \circ Z|_E \not\geq 0$ for every covector Z (of \widetilde{M}') that is 0 at g and $+$ at f . Since Y is a covector containing g in \widetilde{M}' , $Y \setminus \{f\}$ is a covector of \widetilde{M} containing g . By Lemma 7.1.8, $\underline{Y} \setminus \{f\}$ contains a cocircuit (of \widetilde{M}) whose support is of the form $(E \setminus B_0) \cup \{g\}$ for some $B_0 \in \mathcal{B}(M)$. If the containment is proper, then $\underline{Y} \setminus \{f\}$ contains some cocircuit Z_0 of \underline{M} . Since the extension

is generic, the extension Z'_0 of Z_0 in M' contains f . By abusing notation, we consider Z'_0 as a signed cocircuit of M' (hence \widetilde{M}') in which $Z'_0(f) = +$. Now we have a contradiction as $Y \circ Z_0|_E \geq 0$. Therefore $\underline{Y} \setminus \{f\} = (E \setminus B_0) \cup \{g\}$, and it is a cocircuit of \widetilde{M} . We claim that B_0 is the basis of M we want.

The second assertion is immediate as $Y|_{E \cup \{g\}} \geq 0$ (recall that after reorienting f, g , we are working with \mathcal{O}'_+ and $\widetilde{\mathcal{O}}_+$ instead). By Lemma 7.1.8, $B_0 \cup \{f\}$ is a circuit of \underline{M}' . Denote by X the signed circuit of M' supported on $B_0 \cup \{f\}$ such that $X(f) = +$, it remains to show $X \geq 0$. Suppose $X(e) = -$. Let Z_e be the fundamental cocircuit of e with respect to B_0 in \underline{M} , and let Z'_e be its extension in \underline{M}' . Since the extension is generic, $f \in \underline{Z}'_0$, and again we can abuse notation to consider Z'_e as a signed cocircuit of M' (hence \widetilde{M}') in which $Z'_e(f) = +$. From the choice of Z'_e , $\underline{Z}'_e \cap \underline{X} = \{e, f\}$, so $Z'_e(e) = +$ by orthogonality. In particular, $Y \circ Z_e|_E \geq 0$, which is a contradiction. Therefore $B_0 \cup \{f\}$ is a positive circuit of \mathcal{O}'_+ as well. \square

PROOF OF THEOREM 7.1.5: By Proposition 7.1.10 and 7.1.11, the images of $\hat{\beta}_{\sigma, \sigma^*}$ are (σ, σ^*) -compatible. Injectivity follows from Proposition 7.1.10 and the uniqueness assertion of Theorem 7.1.6. Surjectivity follows from Proposition 7.1.10 and the existence assertion of Theorem 7.1.6. \square

As a straightforward corollary, we have the following generalization of Theorem 4.3.6.

Corollary 7.1.12 *The number of (σ, σ^*) -compatible orientations of an oriented matroid M equals the number of bases of M .*

7.2 Simplicial and Cellular Trees

A graph can be thought as a 1-dimensional simplicial complex, and spanning trees are subcomplexes that satisfy certain conditions in terms of algebraic topology. A higher-dimensional analogue of spanning trees was studied by Duval, Klivans and Martin [47]. We

give a brief introduction of their theory, focusing its matroidal side and connections with our work. We will be using basic terminology and theory in algebraic topology, which can be found in standard texts such as [64].

We fix some notations first. Throughout this section, Δ will be a finite, pure cell complex of dimension d . The collection of i -dimensional faces of Δ will be denote by Δ_i , and $\Delta_{\leq i} := \bigsqcup_{j \leq i} \Delta_j$ is the i -skeleton of Δ . The *module of i -chains* $C_i(\Delta; R)$ is the free R -module generated by Δ_i ; denote by $\partial_i : C_i(\Delta; R) \rightarrow C_{i-1}(\Delta; R)$ and $\partial_i^* : C_{i-1}(\Delta; R) \rightarrow C_i(\Delta; R)$ the usual *boundary map* and *coboundary map* in cellular (co)homology, respectively. Finally, given an abelian group G , denote by $\mathbf{T}(G)$ the *torsion subgroup* of G .

Definition 7.2.1 *A subcomplex Υ of Δ is a maximal cellular spanning forest of Δ if*

1. $\Upsilon_{\leq d-1} = \Delta_{\leq d-1}$.
2. $\tilde{H}_d(\Upsilon; \mathbb{Z}) = \{0\}$.
3. $\dim_{\mathbb{Q}} \tilde{H}_{d-1}(\Upsilon; \mathbb{Q}) = \dim_{\mathbb{Q}} \tilde{H}_{d-1}(\Delta; \mathbb{Q})$.

In general, a subcomplex is a cellular spanning forest if it satisfies the first two conditions. If Δ itself is connected, that is, $\dim_{\mathbb{Q}} \tilde{H}_{d-1}(\Delta; \mathbb{Q}) = 0$, then a maximal cellular spanning forest of Δ is called a cellular spanning tree.

Remark. The definition of connectedness here is different from the usual topological sense that $\tilde{H}_0(\Delta; \mathbb{Q}) = \{0\}$. So unlike the case of graphs where one can study maximal spanning forests by studying spanning trees of each connected component, one has to work with spanning forests as inseparable objects.

By fixing a reference orientation for every face of Δ , the boundary map ∂_i can be expressed as a $\Delta_{i-1} \times \Delta_i$ (integer) matrix in which the (F', F) -th entry is the multiplicity of F' on the boundary of F , and the coboundary map ∂_i^* can be expressed as the transpose

of ∂_i . The topological definition above can be translated into a more direct definition using matroidal language.

Proposition 7.2.2 [47, Proposition 2.13] *A collection of facets, together with $\Delta_{\leq d-1}$, form a cellular spanning forest if and only if the collection is an independent set of the matroid $M(\Delta) := M(\partial_d)$. Hence a collection of facets is a maximal cellular spanning forest if and only if it is a basis of $M(\Delta)$.*

By applying Theorem 5.4.2 to $M(\Delta)$, we have the following formal corollary.

Corollary 7.2.3 *Let σ (resp. σ^*) be an arbitrary acyclic circuit (resp. cocircuit) signature of $M(\Delta)$. The number of cellular spanning forests of Δ equals the number of σ -compatible orientations of $M(\Delta)$; the number of maximal cellular spanning forests of Δ equals the number of (σ, σ^*) -compatible orientations of $M(\Delta)$, in which the map $\hat{\beta}_{\sigma, \sigma^*}$ defined in Theorem 5.4.2 provides an explicit bijection.*

We emphasize that while the corollary is presented in matroidal terms, these terms have certain topological intuition. Every element ρ of $M(\Delta)$ corresponds to a d -dimensional cell of Δ , which is a copy of the disk D^d glued to $\Delta_{\leq d-1}$ along boundary, and the orientation of $\rho \in M(\Delta)$ corresponds to the usual topological orientation of the cell $\rho \in \Delta$.

Now we discuss a bit on the topological meaning of the map $\hat{\beta}_{\sigma, \sigma^*}$. The fundamental circuit of a facet ρ with respect to a maximal spanning forest Υ in $M(\Delta)$ corresponds to a non-trivial top homology class of $\Upsilon \cup \rho$, which, roughly speaking, can be thought as a collection of facets enclosing a $(d+1)$ -dimensional “hole”. Such “hole” can be filled in by a $(d+1)$ -dimensional cell, and the orientation of the circuit is induced by the orientation of such phantom cell.

Dually, the matroidal notion of a signed fundamental cocircuit can be interpreted as follows: for any other facet ρ' in the fundamental cocircuit of ρ with respect to Υ , $\Upsilon \cup \rho'$ encloses a $(d+1)$ -dimensional “hole” whose boundary includes both ρ and ρ' , and an

orientation of ρ induces an orientation of ρ' via the orientation of the “hole” and vice versa, this gives a topological way to assign consistent orientation across the cocircuit.

A theme in the theory of cellular forests is that one usually counts the number of maximal spanning forests according to multiplicities. More precisely, the *tree number* $\tau(\Delta)$ is the sum $\sum_{\Upsilon} |\mathbf{T}(\tilde{H}_{d-1}(\Upsilon; \mathbb{Z}))|^2$ taken over all maximal spanning forests Υ of Δ . Tree number was introduced by Gil Kalai [71] to formulate a generalization of Cayley’s formula, namely that the tree number of the d -skeleton of a n -simplex is $n^{\binom{n-2}{d}}$. Tree number is a more well-behaved invariant than the direct count of maximal spanning forests in many senses. As an example, counting the number of maximal spanning forests is #P-Hard [101], while the tree number can be computed efficiently.

We have the following generalization of Laplacians and Jacobians, as well as generalization of the Matrix–Tree theorem, known as the Cellular Matrix–Tree Theorem.

Definition 7.2.4 *The (up-down) Laplacian $L := L_{\Delta}$ is defined as $\partial_d \partial_d^*$. The critical group $K(\Delta)$ of Δ is $\mathbf{T}(\ker_{\mathbb{Z}} \partial_{d-1} / \text{im}_{\mathbb{Z}} L)$.*

The flow lattice \mathcal{F} and the cut lattice \mathcal{C} of M are $\ker_{\mathbb{Z}} \partial_d$ and $\text{im}_{\mathbb{Z}} \partial_d^$, respectively. The cutflow group of Δ is $C_d(\Delta; \mathbb{Z}) / (\mathcal{F} \oplus \mathcal{C})$.*

Theorem 7.2.5 [46, Proposition 3.5, Theorem 8.1] *Pick an arbitrary $\Gamma \subset \Delta_{d-1}$ such that the rows of ∂_d corresponding to $\Delta_{d-1} \setminus \Gamma$ form a basis of the row space (such Γ should be thought as a root of the cellular forest). Let $L_{\Delta \setminus \Gamma}$ be the restriction of L to the faces of $\Delta_{d-1} \setminus \Gamma$. Then*

$$\tau(\Delta) = \frac{|\mathbf{T}(\tilde{H}_{d-1}(\Delta; \mathbb{Z}))|^2}{|\mathbf{T}(\tilde{H}_{d-1}(\Delta, \Gamma; \mathbb{Z}))|^2} \det L_{\Delta \setminus \Gamma}.$$

In terms of critical groups, we have the following equations.

$$\tau(\Delta) = |K(\Delta)| = |C_d(\Delta; \mathbb{Z}) / (\mathcal{F} \oplus \mathcal{C})| |\mathbf{T}(\tilde{H}_{d-1}(\Delta; \mathbb{Z}))|.$$

Because of the extra data in Theorem 7.2.5, namely the cardinality of the homology

groups, $M(\Delta)$ should be considered as an *arithmetic matroid* [40]. There are two natural arithmetic structures on $M(\Delta)$. The first one is induced by the matrix ∂_d , so the *multiplicity* of a maximal spanning forest Υ equals the greatest common divisor of the maximal minors of $\partial_d|_{\Upsilon}$, which is precisely $|\mathbf{T}(\tilde{H}_{d-1}(\Upsilon; \mathbb{Z}))|$ [14], such an arithmetic structure is *realizable*. The second one associates the multiplicity $|\mathbf{T}(\tilde{H}_{d-1}(\Upsilon; \mathbb{Z}))|^2$ to each maximal spanning forest Υ [41].

As an application of our work, we have the following reduction for a case of the sampling question posed at the end of [47, Section 5.4], which is of similar flavor as Proposition 5.8.3. Given a zonotope Z and a generic direction \mathbf{v} (not contained in the affine span of any facet of Z), the *half-open zonotope* (with respect to \mathbf{v}) consists of the points in Z that remain in Z after being shifted by $\epsilon\mathbf{v}$ for some sufficiently small $\epsilon > 0$.

Proposition 7.2.6 *The problem of sampling a random maximal cellular spanning forest, using the distribution proportional to the cardinality of the torsion of its top homology group, is polynomial-time reducible from each of the following two problems:*

1. *uniformly sampling a lattice point from a half-open Z_{∂_d} , and*
2. *uniformly sampling a point from $Z_{\partial_d} \subset \mathbb{R}^{\Delta_{d-1}}$ with respect to the relative volume in the affine span of Z_{∂_d} .*

PROOF: Pick an arbitrary generic shifting vector \mathbf{v} and an arbitrary fine zonotope tiling of Z_{∂_d} induced by some acyclic circuit signature. Each parallelepiped \mathcal{C}_{Υ} corresponds to a maximal spanning forest Υ , and the number of lattice points in the half-open \mathcal{C}_{Υ} (with respect to \mathbf{v}) is equal to $|\mathbf{T}(\tilde{H}_{d-1}(\Upsilon; \mathbb{Z}))|$ [39]. Given a lattice point in the half-open Z , we can find the half-open parallelepiped (thus which maximal spanning forest) it is in using a similar inverse algorithm as described in Section 5.8.1, so we can convert a random lattice point into a random maximal spanning forest as stated in the proposition.

The second part follows from the fact that the relative volume of \mathcal{C}_Υ is equal to the number of lattice points in its half-open version [15, Lemma 9.2], and that the same inverse algorithm converts a random point to a random maximal spanning forest. \square

On the other hand, the second arithmetic structure on $M(\Delta)$ is often non-realizable [82], so finding the geometric meanings of such structure (and their algorithmic consequences) is an interesting problem.

As a final note, almost all discussions in this section become nicer when ∂_d is a totally unimodular matrix. Examples and properties of such complexes are studied in [14, 49].

REFERENCES

- [1] Y. An, M. Baker, G. Kuperberg, and F. Shokrieh, “Canonical Representatives for Divisor Classes on Tropical Curves and the Matrix-tree Theorem,” *Forum Math. Sigma*, vol. 2, e24, 25, 2014.
- [2] M. Arbo and N. Proudfoot, “Hypertoric Varieties and Zonotopal Tilings,” *Int. Math. Res. Not. IMRN*, no. 23, pp. 7268–7301, 2016.
- [3] R. Bacher, P. de la Harpe, and T. Nagnibeda, “The Lattice of Integral Flows and the Lattice of Integral Cuts on a Finite Graph,” *Bull. Soc. Math. France*, vol. 125, no. 2, pp. 167–198, 1997.
- [4] S. Backman, “Riemann-Roch Theory for Graph Orientations,” *Adv. Math.*, vol. 309, pp. 655–691, 2017.
- [5] ———, “Partial Graph Orientations and the Tutte Polynomial,” *Adv. in Appl. Math.*, vol. 94, pp. 103–119, 2018.
- [6] S. Backman, M. Baker, and C. H. Yuen, “Geometric Bijections for Regular Matroids, Zonotopes, and Ehrhart Theory,” Preprint available at <https://arxiv.org/abs/1701.01051>, 2017.
- [7] S. Backman and S. Hopkins, “Fourorientations and the Tutte Polynomial,” *Res. Math. Sci.*, vol. 4, p. 18, 2017.
- [8] P. Bak, C. Tang, and K. Wiesenfeld, “Self-organized Criticality: An Explanation of the $1/f$ Noise,” *Phys. Rev. Lett.*, vol. 59, pp. 381–384, 4 1987.
- [9] M. Baker, “Specialization of Linear Systems from Curves to Graphs,” *Algebra & Number Theory*, vol. 2, no. 6, pp. 613–653, 2008, With an appendix by Brian Conrad.
- [10] M. Baker and X. Faber, “Metric Properties of the Tropical Abel–Jacobi Map,” *J. Algebraic Combin.*, vol. 33, no. 3, pp. 349–381, 2011.
- [11] M. Baker and S. Norine, “Riemann–Roch and Abel–Jacobi Theory on a Finite Graph,” *Adv. Math.*, vol. 215, no. 2, pp. 766–788, 2007.
- [12] M. Baker and F. Shokrieh, “Chip-firing Games, Potential Theory on Graphs, and Spanning Trees,” *J. Combin. Theory Ser. A*, vol. 120, no. 1, pp. 164–182, 2013.

- [13] M. Baker and Y. Wang, “The Bernardi Process and Torsor Structures on Spanning Trees,” *Int. Math. Res. Not. IMRN*, rnx037, 2017.
- [14] M. Beck, F. Breuer, L. Godkin, and J. L. Martin, “Enumerating Colorings, Tensions and Flows in Cell Complexes,” *J. Combin. Theory Ser. A*, vol. 122, pp. 82–106, 2014.
- [15] M. Beck and S. Robins, *Computing the Continuous Discretely*, Second, ser. Undergraduate Texts in Mathematics. Springer, New York, 2015, pp. xx+285, Integer-point enumeration in polyhedra, With illustrations by David Austin, ISBN: 978-1-4939-2968-9; 978-1-4939-2969-6.
- [16] G. Benkart, C. Klivans, and V. Reiner, “Chip Firing on Dynkin Diagrams and McKay Quivers,” *Math. Z.*, 2018.
- [17] B. Benson, D. Chakrabarty, and P. Tetali, “ G -parking Functions, Acyclic Orientations and Spanning Trees,” *Discrete Math.*, vol. 310, no. 8, pp. 1340–1353, 2010.
- [18] M. Bern and D. Eppstein, “Optimization over Zonotopes and Training Support Vector Machines,” in *Algorithms and Data structures (Providence, RI, 2001)*, ser. Lecture Notes in Comput. Sci. Vol. 2125, Springer, Berlin, 2001, pp. 111–121.
- [19] O. Bernardi, “A Characterization of the Tutte Polynomial via Combinatorial Embeddings,” *Ann. Comb.*, vol. 12, no. 2, pp. 139–153, 2008.
- [20] ———, “Tutte Polynomial, Subgraphs, Orientations and Sandpile Model: New Connections via Embeddings,” *Electron. J. Combin.*, vol. 15, no. 1, 2008.
- [21] N. L. Biggs, “Chip-firing and the Critical Group of a Graph,” *J. Algebraic Combin.*, vol. 9, no. 1, pp. 25–45, 1999.
- [22] L. J. Billera and B. Sturmfels, “Fiber Polytopes,” *Ann. of Math. (2)*, vol. 135, no. 3, pp. 527–549, 1992.
- [23] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, and G. M. Ziegler, *Oriented Matroids*, Second, ser. Encyclopedia of Mathematics and its Applications. Cambridge University Press, Cambridge, 1999, vol. 46, pp. xii+548, ISBN: 0-521-77750-X.
- [24] A. Björner, L. Lovász, and P. W. Shor, “Chip-firing Games on Graphs,” *European J. Combin.*, vol. 12, no. 4, pp. 283–291, 1991.
- [25] R. G. Bland and M. Las Vergnas, “Orientability of Matroids,” *J. Combinatorial Theory Ser. B.*, vol. 24, no. 1, pp. 94–123, 1978.

- [26] J. Bohne, *Eine Kombinatorische Analyse Zonotopaler Raumaufteilungen*, Ph.D. thesis, Universität Bielefeld., 1992.
- [27] B. Bollobás, *Modern Graph Theory*, ser. Graduate Texts in Mathematics. Springer-Verlag, New York, 1998, vol. 184, pp. xiv+394, ISBN: 0-387-98488-7.
- [28] T. Brylawski and J. Oxley, “The Tutte Polynomial and its Applications,” in *Matroid Applications*, ser. Encyclopedia Math. Appl. Vol. 40, Cambridge Univ. Press, Cambridge, 1992, pp. 123–225.
- [29] L. Caporaso and K. Christ, “Combinatorics of Compactified Universal Jacobians,” Preprint available at <https://arxiv.org/abs/1801.04098>, 2018.
- [30] M. Chan, T. Church, and J. A. Grochow, “Rotor-routing and Spanning Trees on Planar graphs,” *Int. Math. Res. Not. IMRN*, no. 11, pp. 3225–3244, 2015.
- [31] M. Chan, D. Glass, M. Macauley, D. Perkinson, C. Werner, and Q. Yang, “Sand-piles, Spanning Trees, and Plane Duality,” *SIAM J. Discrete Math.*, vol. 29, no. 1, pp. 461–471, 2015.
- [32] D. Chebikin and P. Pylyavskyy, “A Family of Bijections Between G -parking Functions and Spanning Trees,” *J. Combin. Theory Ser. A*, vol. 110, no. 1, pp. 31–41, 2005.
- [33] F. Cools, J. Draisma, S. Payne, and E. Robeva, “A Tropical Proof of the Brill-Noether Theorem,” *Adv. Math.*, vol. 230, no. 2, pp. 759–776, 2012.
- [34] R. Cori and Y. Le Borgne, “The Sand-pile Model and Tutte Polynomials,” *Adv. in Appl. Math.*, vol. 30, no. 1-2, pp. 44–52, 2003, Formal power series and combinatorics (Scottsdale, AZ, 2001).
- [35] J. Courtiel, “A General Notion of Activity for the Tutte Polynomial,” Preprint available at <https://arxiv.org/abs/1412.2081>, 2014.
- [36] D. A. Cox, J. Little, and D. O’Shea, *Using Algebraic Geometry*, Second, ser. Graduate Texts in Mathematics. Springer, New York, 2005, vol. 185, pp. xii+572, ISBN: 0-387-20706-6.
- [37] H. S. M. Coxeter, “The Classification of Zonohedra by Means of Projective Diagrams,” *J. Math. Pures Appl. (9)*, vol. 41, pp. 137–156, 1962.
- [38] ———, *Regular Polytopes*, ser. Second edition. The Macmillan Co., New York; Collier-Macmillan Ltd., London, 1963, xx+321 pp. (8 plates).

- [39] M. D’Adderio and L. Moci, “Ehrhart Polynomial and Arithmetic Tutte polynomial,” *European J. Combin.*, vol. 33, no. 7, pp. 1479–1483, 2012.
- [40] ———, “Arithmetic Matroids, the Tutte Polynomial and Toric Arrangements,” *Adv. Math.*, vol. 232, pp. 335–367, 2013.
- [41] E. Delucchi and L. Moci, “Products of Arithmetic Matroids and Quasipolynomial Invariants of CW-complexes,” *J. Combin. Theory Ser. A*, vol. 157, pp. 28–40, 2018.
- [42] D. Dhar, “Self-organized Critical State of Sandpile Automaton Models,” *Phys. Rev. Lett.*, vol. 64, no. 14, pp. 1613–1616, 1990.
- [43] ———, “Theoretical Studies of Self-organized Criticality,” *Phys. A*, vol. 369, no. 1, pp. 29–70, 2006.
- [44] R. Diestel, *Graph Theory*, Third, ser. Graduate Texts in Mathematics. Springer-Verlag, Berlin, 2005, vol. 173, pp. xvi+411, ISBN: 978-3-540-26182-7; 3-540-26182-6; 978-3-540-26183-4.
- [45] D. Durfee, J. Peebles, R. Peng, and A. B. Rao, “Determinant-preserving Sparsification of SDDM Matrices with Applications to Counting and Sampling Spanning Trees,” in *58th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2017*, IEEE Computer Soc., Los Alamitos, CA, 2017, pp. 926–937.
- [46] A. Duval, C. Klivans, and J. Martin, “Cuts and Flows of Cell Complexes,” *J. Algebraic Combin.*, vol. 41, no. 4, pp. 969–999, 2015.
- [47] ———, “Simplicial and Cellular Trees,” in *Recent Trends in Combinatorics*, Springer International Publishing, 2016, pp. 713–752, ISBN: 978-3-319-24298-9.
- [48] E. Ehrhart, “Sur les Polyèdres Homothétiques Bordés à n Dimensions,” *C. R. Acad. Sci. Paris*, vol. 254, pp. 988–990, 1962.
- [49] J. Erickson, “Combinatorial Optimization of Cycles and Bases,” in *Advances in Applied and Computational Topology*, ser. Proc. Sympos. Appl. Math. Vol. 70, Amer. Math. Soc., Providence, RI, 2012, pp. 195–228.
- [50] E. S. Fedorov, “Elemente der Gestaltenlehre,” *Z. Kristallogr.*, vol. 21, pp. 671–694, 1893.
- [51] J. Folkman and J. Lawrence, “Oriented Matroids,” *J. Combin. Theory Ser. B*, vol. 25, no. 2, pp. 199–236, 1978.
- [52] A. Frieze, N. Goyal, L. Rademacher, and S. Vempala, “Expanders via Random Spanning Trees,” *SIAM J. Comput.*, vol. 43, no. 2, pp. 497–513, 2014.

- [53] A. Gathmann and M. Kerber, “A Riemann-Roch Theorem in Tropical Geometry,” *Math. Z.*, vol. 259, no. 1, pp. 217–230, 2008.
- [54] I. M. Gelfand, M. M. Kapranov, and A. V. Zelevinsky, *Discriminants, Resultants and Multidimensional Determinants*, ser. Modern Birkhäuser Classics. Birkhäuser Boston, Inc., 2008, pp. x+523, Reprint of the 1994 edition, ISBN: 978-0-8176-4770-4.
- [55] I. M. Gessel and B. E. Sagan, “The Tutte Polynomial of a Graph, Depth-First Search, and Simplicial Complex Partitions,” *Electron. J. Combin.*, vol. 3, no. 2, Research Paper 9, approx. 36, 1996, The Foata Festschrift.
- [56] E. Gioan, *Correspondance Naturelle Entre Bases et Réorientations des Matroïdes Orientés*, Ph.D. thesis, University of Bordeaux 1, 2002. Available at <http://www.lirmm.fr/~gioan>, 2002.
- [57] —, “Enumerating Degree Sequences in Digraphs and a Cycle–cocycle Reversing System,” *European Journal of Combinatorics*, vol. 28, no. 4, pp. 1351–1366, 2007.
- [58] —, “Circuit-cocircuit Reversing Systems in Regular Matroids,” *Ann. Comb.*, vol. 12, no. 2, pp. 171–182, 2008.
- [59] C. Greene and T. Zaslavsky, “On the Interpretation of Whitney Numbers through Arrangements of Hyperplanes, Zonotopes, non-Radon Partitions, and Orientations of Graphs,” *Trans. Amer. Math. Soc.*, vol. 280, no. 1, pp. 97–126, 1983.
- [60] P. Griffiths and J. Harris, “On the Variety of Special Linear Systems on a General Algebraic Curve,” *Duke Math. J.*, vol. 47, no. 1, pp. 233–272, 1980.
- [61] J. L. Gross and T. W. Tucker, *Topological Graph Theory*, ser. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons, Inc., New York, 1987, pp. xvi+351, A Wiley-Interscience Publication, ISBN: 0-471-04926-3.
- [62] J. Guzmán and C. Klivans, “Chip Firing on General Invertible Matrices,” *SIAM J. Discrete Math.*, vol. 30, no. 2, pp. 1115–1127, 2016.
- [63] C. Haase, G. Musiker, and J. Yu, “Linear Systems on Tropical Curves,” *Math. Z.*, vol. 270, no. 3-4, pp. 1111–1140, 2012.
- [64] A. Hatcher, *Algebraic Topology*. Cambridge University Press, Cambridge, 2002, pp. xii+544, ISBN: 0-521-79160-X; 0-521-79540-0.
- [65] A. E. Holroyd, L. Levine, K. Mészáros, Y. Peres, J. Propp, and D. B. Wilson, “Chip-firing and Rotor-routing on Directed Graphs,” in *In and Out of Equilibrium. 2*, ser. Progr. Probab. Vol. 60, Birkhäuser, Basel, 2008, pp. 331–364.

- [66] O. Holtz and A. Ron, “Zonotopal Algebra,” *Adv. Math.*, vol. 227, no. 2, pp. 847–894, 2011.
- [67] S. Hopkins, “Problems from the AIM Chip-Firing Workshop,” Available at http://web.mit.edu/~shopkins/docs/aim_chip-firing_problems.pdf, 2013.
- [68] ———, “Open Problems from the CMO-BIRS Workshop on Sandpile Groups,” Available at http://web.mit.edu/shopkins/docs/birs_sandpile_problems.pdf, 2015.
- [69] B. Huber, J. Rambau, and F. Santos, “The Cayley Trick, Lifting Subdivisions and the Bohne-Dress Theorem on Zonotopal Tilings,” *J. Eur. Math. Soc. (JEMS)*, vol. 2, no. 2, pp. 179–198, 2000.
- [70] D. C. Jerison, L. Levine, and J. Pike, “Mixing Time and Eigenvalues of the Abelian Sandpile Markov Chain,” Preprint available at <https://arxiv.org/abs/1511.00666>, 2015.
- [71] G. Kalai, “Enumeration of \mathbb{Q} -acyclic Simplicial Complexes,” *Israel J. Math.*, vol. 45, no. 4, pp. 337–351, 1983.
- [72] E. Kaltofen and G. Villard, “On the Complexity of Computing Determinants,” *Comput. Complexity*, vol. 13, no. 3-4, pp. 91–130, 2004.
- [73] R. Kannan and S. Vempala, “Sampling Lattice Points,” in *STOC '97 (El Paso, TX)*, ACM, New York, 1999, pp. 696–700.
- [74] J. Kass and N. Pagani, *Private Correspondence*.
- [75] R. Kenyon and P. Winkler, “Branched polymers,” *Amer. Math. Monthly*, vol. 116, no. 7, pp. 612–628, 2009.
- [76] G. Kirchhoff, “Ueber die Auflösung der Gleichungen, auf Welche man bei der Untersuchung der Linearen Vertheilung Galvanischer Ströme geführt wird,” *Annalen der Physik*, vol. 148, pp. 497–508, 1847.
- [77] J. P. S. Kung, “Convolution–Multiplication Identities for Tutte Polynomials of Graphs and Matroids,” *J. Combin. Theory Ser. B*, vol. 100, no. 6, pp. 617–624, 2010.
- [78] M. Las Vergnas, “Convexity in Oriented Matroids,” *J. Combin. Theory Ser. B*, vol. 29, no. 2, pp. 231–243, 1980.

- [79] ———, “The Tutte Polynomial of a Morphism of Matroids. II. Activities of Orientations,” in *Progress in Graph Theory (Waterloo, Ont., 1982)*, Academic Press, Toronto, ON, 1984, pp. 367–380.
- [80] E. L. Lawler, “Matroid Intersection Algorithms,” *Math. Programming*, vol. 9, no. 1, pp. 31–56, 1975.
- [81] J. M. Lee, *Introduction to Smooth Manifolds*, ser. Graduate Texts in Mathematics. Springer-Verlag, New York, 2003, vol. 218, pp. xviii+628, ISBN: 0-387-95495-3.
- [82] M. Lenz, “On Powers of Plücker Coordinates and Representability of Arithmetic Matroids,” Preprint available at <https://arxiv.org/abs/1703.10520>, 2017.
- [83] D. J. Lorenzini, “On the Group of Components of a Néron Model,” *J. Reine Angew. Math.*, vol. 445, pp. 109–160, 1993.
- [84] A. Madry, “Navigating Central Path with Electrical Flows: From Flows to Matchings, and Back,” in *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, 2013, pp. 253–262.
- [85] P. McMullen, “On Zonotopes,” *Trans. Amer. Math. Soc.*, vol. 159, pp. 91–109, 1971.
- [86] C. Merino López, “Chip Firing and the Tutte Polynomial,” *Ann. Comb.*, vol. 1, no. 3, pp. 253–259, 1997.
- [87] ———, *Matroids, the Tutte Polynomial and the Chip Firing Game*, PhD thesis, University of Oxford., 1999.
- [88] G. Mikhalkin and I. Zharkov, “Tropical Curves, Their Jacobians and Theta Functions,” in *Curves and Abelian Varieties*, ser. Contemp. Math. Vol. 465, Amer. Math. Soc., Providence, RI, 2008, pp. 203–230.
- [89] G. J. Minty, “On the Axiomatic Foundations of the Theories of Directed Linear Graphs, Electrical Networks and Network-programming,” *J. Math. Mech.*, vol. 15, pp. 485–520, 1966.
- [90] F. Mohammadi and F. Shokrieh, “Divisors on Graphs, Binomial and Monomial Ideals, and Cellular Resolutions,” *Math. Z.*, vol. 283, no. 1-2, pp. 59–102, 2016.
- [91] R. Nickel, *Flows and Colorings in Oriented Matroids*, Ph.D. thesis, Fernuniversität Hagen. Available at <https://d-nb.info/1031144978/34>, 2012.

- [92] J. Oxley, *Matroid Theory*, Second, ser. Oxford Graduate Texts in Mathematics. Oxford University Press, Oxford, 2011, vol. 21, pp. xiv+684, ISBN: 978-0-19-960339-8.
- [93] D. Perkinson, J. Perlman, and J. Wilmes, “Primer for the Algebraic Geometry of Sandpiles,” in *Tropical and Non-Archimedean Geometry*, ser. Contemp. Math. Vol. 605, Amer. Math. Soc., Providence, RI, 2013, pp. 211–256.
- [94] D. Perkinson, Q. Yang, and K. Yu, “ G -parking Functions and Tree Inversions,” *Combinatorica*, vol. 37, no. 2, pp. 269–282, 2017.
- [95] A. Postnikov and B. Shapiro, “Trees, Parking Functions, Syzygies, and Deformations of Monomial Ideals,” *Trans. Amer. Math. Soc.*, vol. 356, no. 8, pp. 3109–3142, 2004.
- [96] J. Richter-Gebert and G. M. Ziegler, “Zonotopal Tilings and the Bohne-Dress Theorem,” in *Jerusalem Combinatorics ’93*, ser. Contemp. Math. Vol. 178, Amer. Math. Soc., Providence, RI, 1994, pp. 211–232.
- [97] F. Santos, “Triangulations of Oriented Matroids,” *Mem. Amer. Math. Soc.*, vol. 156, no. 741, pp. viii+80, 2002.
- [98] A. Schrijver, *Theory of Linear and Integer Programming*. Wiley, 1998.
- [99] J. Shen, *Break Divisors and Compactified Jacobians*, Ph.D. thesis, Yale University., 2017.
- [100] C. T. Simpson, “Moduli of Representations of the Fundamental Group of a Smooth Projective Variety. I,” *Inst. Hautes Études Sci. Publ. Math.*, no. 79, pp. 47–129, 1994.
- [101] M. Snook, “Counting Bases of Representable Matroids,” *Electron. J. Combin.*, vol. 19, no. 4, Paper 41, 11, 2012.
- [102] R. P. Stanley, “A Zonotope Associated with Graphical Degree Sequences,” in *Applied Geometry and Discrete Mathematics*, ser. DIMACS Ser. Discrete Math. Theoret. Comput. Sci. Vol. 4, Amer. Math. Soc., Providence, RI, 1991, pp. 555–570.
- [103] B. Sturmfels, “On the Newton Polytope of the Resultant,” *J. Algebraic Combin.*, vol. 3, no. 2, pp. 207–236, 1994.
- [104] ———, *Gröbner Bases and Convex Polytopes*, ser. University Lecture Series. American Mathematical Society, Providence, RI, 1996, vol. 8, pp. xii+162, ISBN: 0-8218-0487-1.

- [105] Y. Su and D. G. Wagner, “The Lattice of Integer Flows of a Regular Matroid,” *J. Combin. Theory Ser. B*, vol. 100, no. 6, pp. 691–703, 2010.
- [106] W. T. Tutte, “A Contribution to the Theory of Chromatic Polynomials,” *Canadian J. Math.*, vol. 6, pp. 80–91, 1954.
- [107] F. Vallentin, “A Note on Space Tiling Zonotopes,” Preprint available at <https://arxiv.org/abs/math/0402053>, 2004.
- [108] G. Voronoi, “Nouvelles Applications des Paramètres Continus à théorie des Formes Quadratiques. Deuxième Mémoire. Recherches sur les Paralléloèdres Primitifs,” *J. Reine Angew. Math.*, vol. 136, pp. 67–182, 1909.
- [109] N. White, Ed., *Combinatorial Geometries*, ser. Encyclopedia of Mathematics and its Applications. Cambridge University Press, Cambridge, 1987, vol. 29, pp. xii+212, ISBN: 0-521-33339-3.
- [110] H. Whitney, “On the Abstract Properties of Linear Dependence,” *Amer. J. Math.*, vol. 57, no. 3, pp. 509–533, 1935.
- [111] D. B. Wilson, “Generating Random Spanning Trees More Quickly than the Cover Time,” in *Proceedings of the Twenty-eighth Annual ACM Symposium on the Theory of Computing (Philadelphia, PA, 1996)*, ACM, New York, 1996, pp. 296–303.
- [112] M. M. Wood, “The Distribution of Sandpile Groups of Random Graphs,” *J. Amer. Math. Soc.*, vol. 30, no. 4, pp. 915–958, 2017.
- [113] C. H. Yan, “Parking Functions,” in *Handbook of Enumerative Combinatorics*, ser. Discrete Math. Appl. (Boca Raton), CRC Press, Boca Raton, FL, 2015, pp. 835–893.
- [114] C. H. Yuen, “Geometric Bijections Between Spanning Trees and Break Divisors,” *J. Combin. Theory Ser. A*, vol. 152, no. 1, pp. 159–189, 2017.
- [115] G. M. Ziegler, *Lectures on Polytopes*, ser. Graduate Texts in Mathematics. Springer-Verlag, New York, 1995, vol. 152, pp. x+370, ISBN: 0-387-94365-X.