

ROOT SYSTEM CHIP-FIRING I: INTERVAL-FIRING

PAVEL GALASHIN, SAM HOPKINS, THOMAS MCCONVILLE,
AND ALEXANDER POSTNIKOV

ABSTRACT. Jim Propp recently introduced a variant of chip-firing on a line where the chips are given distinct integer labels. Hopkins, McConville, and Propp showed that this process is confluent from some (but not all) initial configurations of chips. We recast their set-up in terms of root systems: labeled chip-firing can be seen as a *root-firing* process which allows the moves $\lambda \rightarrow \lambda + \alpha$ for $\alpha \in \Phi^+$ whenever $\langle \lambda, \alpha^\vee \rangle = 0$, where Φ^+ is the set of positive roots of a root system of Type A and λ is a weight of this root system. We are thus motivated to study the exact same root-firing process for an arbitrary root system. Actually, this *central root-firing* process is the subject of a sequel to this paper. In the present paper, we instead study the *interval root-firing* processes determined by $\lambda \rightarrow \lambda + \alpha$ for $\alpha \in \Phi^+$ whenever $\langle \lambda, \alpha^\vee \rangle \in [-k - 1, k - 1]$ or $\langle \lambda, \alpha^\vee \rangle \in [-k, k - 1]$, for any $k \geq 0$. We prove that these interval-firing processes are always confluent, from any initial weight. We also show that there is a natural way to consistently label the stable points of these interval-firing processes across all values of k so that the number of weights with given stabilization is a polynomial in k . We conjecture that these *Ehrhart-like polynomials* have nonnegative integer coefficients.

CONTENTS

1. Introduction	2
Part 1. Confluence of symmetric and truncated interval-firing	7
2. Background on root systems	7
3. Background on binary relations and confluence	11
4. Definition of interval-firing	12
5. Symmetries of interval-firing processes	19
6. Sinks of symmetric interval-firing and the map η	21
7. Traverse lengths of permutohedra	24
8. The permutohedron non-escaping lemma	28
9. Confluence of symmetric interval-firing	30
10. Full-dimensional components, saturated components, and Cartan matrix chip-firing as a limit	31
11. Confluence of truncated interval-firing	33

Date: April 14, 2018.

2010 Mathematics Subject Classification. 17B22; 52B20; 05C57.

Key words and phrases. Chip-firing; Abelian Sandpile Model; root systems; confluence; permutohedra; Ehrhart polynomials; zonotopes.

Part 2. Ehrhart-like polynomials	37
12. Ehrhart-like polynomials: introduction	37
13. Symmetric Ehrhart-like polynomials	38
14. Cubical subcomplexes	42
15. How interval-firing components decompose	43
16. Truncated Ehrhart-like polynomials	48
17. Iterative descriptions of the stabilization	50
References	51

1. INTRODUCTION

The Abelian Sandpile Model (ASM) is a discrete dynamical system that takes place on a graph. The states of this system are configurations of grains of sand on the vertices of the graph. A vertex with at least as many grains of sand as its neighbors is said to be *unstable*. Any unstable vertex may *topple*, sending one grain of sand to each of its neighbors. The sequence of topplings may continue forever, or it may terminate at a *stable* configuration, where every vertex is stable. The ASM was introduced (in the special case of the two-dimensional square lattice) by the physicists Bak, Tang, and Wiesenfeld [BTW87] as a simple model of self-organized criticality; much of the general, graphical theory was subsequently developed by Dhar [Dha90, Dha99]. The ASM is by now studied in many parts of both physics and pure mathematics: for instance, following the seminal work of Baker and Norine [BN07], it is known that this model is intimately related to tropical algebraic geometry (specifically, divisor theory for tropical curves [GK08, MZ08]); meanwhile, the ASM is studied by probabilists because of its remarkable scaling-limit behavior [PS13, LPS16]; and there are also interesting complexity-theoretic questions related to the ASM, such as, what is the complexity of determining whether a given configuration stabilizes [KT15, FL16]. For more on sandpiles, consult the short survey article [LP10] or the upcoming textbook [CP18].

Independently of its introduction in the statistical mechanics community, the same model was defined and studied from a combinatorial perspective by Björner, Lovász, and Shor [BLS91] under the name of *chip-firing*.¹ Instead of grains of sand, we imagine that chips are placed on the vertices of a graph; the operation of an unstable vertex sending one chip to each of its neighbors is now called *firing* that vertex. One fundamental result of Björner-Lovász-Shor is that, from any initial chip configuration, either the chip-firing process always goes on forever, or it terminates at a stable configuration that does not depend on the choice of which vertices were fired. This is a *confluence* result: it says that (in the case of termination) the divergent paths in the chip-firing process must come together eventually. This confluence property is the essential property which serves as the basis of all further study of the chip-firing process; it explains the adjective “Abelian” in “Abelian Sandpile Model.”

¹It is also worth mentioning that essentially the same model was studied even earlier, in the context of math pedagogy, by Engel [Eng75, Eng76] under the name of the *probabilistic abacus*.


A closely related chip-firing process to the one studied by Björner-Lovász-Shor is where a distinguished vertex is chosen to be the *sink*. The sink will never become unstable and is allowed to accumulate any number of chips; hence, any initial chip configuration will eventually stabilize to a unique stable configuration. This model was studied for instance by Biggs [Big99] and by Dhar [Dha90, Dha99]. Chip-firing with a sink has been generalized to several other contexts beyond graphs. One of the most straightforward but also nicest such generalizations is what is called *M-matrix chip-firing* (see e.g. [Gab93, GK15], or [PS04, §13]). Rather than a graph, we take as input an integer matrix $\mathbf{C} = (C_{ij}) \in \mathbb{Z}^{n \times n}$. The states are vectors $c = (c_1, c_2, \dots, c_n) \in \mathbb{Z}^n$, and a firing move replaces a state c with $c - \mathbf{C}^t e_i$ whenever $c_i \geq C_{ii}$ for $i = 1, \dots, n$. (Here e_1, \dots, e_n are the standard basis vectors of \mathbb{Z}^n ; i.e., $c - \mathbf{C}^t e_i$ is c minus the i th row of \mathbf{C} .) This firing move is denoted $c \rightarrow c - \mathbf{C}^t e_i$. Setting \mathbf{C} to be the reduced Laplacian of a graph (including possibly a directed graph, as in [BL92]) recovers chip-firing with a sink. But in fact \mathbf{C} does not need to be a reduced Laplacian of any graph for confluence to hold in this setting: the condition required to guarantee confluence (and termination), as first established by Gabrielov [Gab93], is that \mathbf{C} be an M-matrix.

We will discuss M-matrix chip-firing, and its relation to our present research, in more detail later (see §10). But now let us explain the direct motivation for our work, namely, *labeled* chip-firing.

Björner, Lovász, and Shor were motivated to introduce the chip-firing process for arbitrary graphs by papers of Spencer [Spe86] and Anderson et al. [ALS⁺89] which studied the special case of chip-firing on a line. Jim Propp recently introduced a version of labeled chip-firing on a line that generalizes this original case. In ordinary chip-firing, the chips are all indistinguishable. But the states of the labeled chip-firing process are configurations of N *distinguishable* chips with integer labels $1, 2, \dots, N$ on the infinite path graph \mathbb{Z} . The firing moves consist of choosing two chips that occupy the same vertex and moving the chip with the lesser label one vertex to the right and the chip with the greater label one vertex to the left. Propp conjectured that if one starts with an even number of chips at the origin, this labeled chip-firing process is confluent and in particular the chips always end up in sorted order. Propp's conjecture was recently proved by Hopkins, McConville, and Propp [HMP17]. Note crucially that confluence does not hold for labeled chip-firing if the initial number of chips at the origin is odd (e.g., three). Hence, compared to all the other models of chip-firing discussed above (for which confluence holds locally and follows from Newman's *diamond lemma* [New42]), confluence is a much subtler property for labeled chip-firing.

The crucial observation that motivated our present research is that we can generalize Propp's labeled chip-firing to "other types," as follows. For any configuration of N labeled chips on the line, if we define the vector $c := (c_1, c_2, \dots, c_N) \in \mathbb{Z}^N$ by setting c_i to be the position of the chip with label i , then for $i < j$ we are allowed to fire chips with label i and j in this configuration as long as c is orthogonal to $e_i - e_j$; and doing so replaces the vector c by $c + (e_i - e_j)$. Note that the vectors $e_i - e_j$ for $1 \leq i < j \leq N$ are exactly the positive roots Φ^+ of the root system Φ of Type A_{N-1} .

So there is a natural candidate for a generalization of Propp's labeled chip-firing to arbitrary (crystallographic) root systems: let Φ be any root system living in some

FIGURE 1. The $k = 1$ symmetric interval-firing process for $\Phi = A_2$.

Euclidean vector space V ; then for a vector $v \in V$ and a positive root $\alpha \in \Phi^+$, we allow the firing move $v \rightarrow v + \alpha$ whenever v is orthogonal to α . We call this process *central root-firing* (or just *central-firing* for short) because we allow a firing move whenever our point v lies on a certain central hyperplane arrangement (namely, the Coxeter arrangement of Φ).

Central-firing is actually the subject of our sequel paper [GHMP17].

In the present paper we instead study two “affine” deformations of central-firing. Let us explain what these deformations look like. First of all, it turns out to be best to interpret the condition “whenever v is orthogonal to α ” as “whenever $\langle v, \alpha^\vee \rangle = 0$,” where $\langle \cdot, \cdot \rangle$ is the standard inner product on V and α^\vee is the *coroot* associated to α . Also, rather than consider all vectors $v \in V$ to be the states of our system, it is better to restrict to a discrete setting where the states are *weights* $\lambda \in P$, where P is the *weight lattice* of Φ (this is akin to only allowing vectors $c \in \mathbb{Z}^N$ above). The central-firing moves thus become

$$\lambda \rightarrow \lambda + \alpha \text{ whenever } \langle \lambda, \alpha^\vee \rangle = 0 \text{ for } \lambda \in P, \alpha \in \Phi^+.$$

The deformations of central-firing we consider involve changing the values of $\langle \lambda, \alpha^\vee \rangle$ at which we allow the firing move $\lambda \rightarrow \lambda + \alpha$ to be some wider interval. In fact, we study two very particular families of intervals. For $k \in \mathbb{Z}_{\geq 0}$, the *symmetric interval root-firing process* is the binary relation $\xrightarrow{\text{sym},k}$ on P defined by

$$\lambda \xrightarrow{\text{sym},k} \lambda + \alpha, \text{ for } \lambda \in P \text{ and } \alpha \in \Phi^+ \text{ with } \langle \lambda, \alpha^\vee \rangle + 1 \in \{-k, -k+1, \dots, k\}$$

and the *truncated interval root-firing process* is the relation $\xrightarrow{\text{tr},k}$ on P defined by

$$\lambda \xrightarrow{\text{tr},k} \lambda + \alpha, \text{ for } \lambda \in P \text{ and } \alpha \in \Phi^+ \text{ with } \langle \lambda, \alpha^\vee \rangle + 1 \in \{-k+1, -k+2, \dots, k\}.$$

We refer to these as *interval-firing* processes for short.

As mentioned, the central-firing process may or may not be confluent, depending on the initial weight we start at (e.g., our comment about three labeled chips above says

that the central-firing process is not confluent from the origin for the root system of Type A_2). The first main result of the present paper is the following, which we prove in §9 and §11.

Theorem 1.1. *For any $k \geq 0$, both the symmetric and truncated interval-firing processes are confluent from all initial weights.*

For example, Figure 1 depicts the $k = 1$ symmetric interval-firing process for $\Phi = A_2$: the edges of this graph correspond to firing moves; that this process is confluent means that all paths starting from a given vertex must terminate at the same final vertex. For more such pictures, see Example 4.1.

We call these processes *interval-firing* processes because they allow firing a root from a weight when the inner product of that weight with the corresponding coroot is in some fixed interval. Alternately, we could say that the firing moves are allowed when our weight belongs to a certain affine hyperplane arrangement whose hyperplanes are orthogonal translates of the Coxeter arrangement hyperplanes; this is precisely the sense in which these processes are “affine.” The *symmetric* process is so called because the symmetric closure of the relation $\xrightarrow{\text{sym},k}$ is invariant under the action of the Weyl group. The *truncated* process is so-called because the interval defining it is truncated by one element on the left compared to the symmetric process.

Note that these processes are not truly “deformations” of central-firing in the sense that we cannot recover central-firing by specializing k . But observe that the $k = 0$ case of symmetric interval-firing has the firing moves

$$\lambda \xrightarrow{\text{sym},0} \lambda + \alpha \text{ whenever } \langle \lambda, \alpha^\vee \rangle = -1 \text{ for } \lambda \in P, \alpha \in \Phi^+$$

and the $k = 1$ case of truncated interval-firing has the firing moves

$$\lambda \xrightarrow{\text{tr},1} \lambda + \alpha \text{ whenever } \langle \lambda, \alpha^\vee \rangle \in \{-1, 0\} \text{ for } \lambda \in P, \alpha \in \Phi^+.$$

So these two interval-firing processes are actually very “close” to central-firing, and suggest that central-firing (in particular, labeled chip-firing) is somehow right on the “cusp” of confluence. Hence, it is not surprising that some of the tools we develop in the present paper are applied to the study of central-firing in the sequel paper [GHMP17]. We also note that these interval-firing processes themselves have a direct chip-firing interpretation in Type A; see Remark 4.2 for more details.

Moreover, we contend that these interval-firing processes are interesting not just because of their connection to central-firing (and hence labeled chip-firing), but also because of their remarkable geometric structure. To get a sense of this geometric structure, the reader is encouraged to look at the depictions of these interval-firing processes for the irreducible rank 2 root systems in Example 4.1. As we will show, the symmetric and truncated interval-firing processes are closely related to *permutohedra*, and indeed we will mostly investigate these processes from the perspective of convex, polytopal geometry. For example, a key ingredient in our proof of confluence is an exact formula for *traverse lengths of root strings* in permutohedra.

The most striking geometric objects that come out of our investigation of interval-firing are certain “Ehrhart-like” polynomials that count the number of weights with

given stabilization as we vary our parameter k . To make sense of “with given stabilization,” first we show that there is a consistent way to label the stable points of the symmetric and truncated interval-firing processes across all values of k : these stable points are (a subset of) $\eta^k(\lambda)$ for $\lambda \in P$, where $\eta: P \rightarrow P$ is a certain piecewise-linear “dilation” map depicted in Figure 8. Then we ask: for $\lambda \in P$, how many weights stabilize to $\eta^k(\lambda)$, as a function of k ? Let us denote by $L_\lambda^{\text{sym}}(k)$ (resp., $L_\lambda^{\text{tr}}(k)$) the number of weights $\mu \in P$ that $\xrightarrow{\text{sym},k}$ -stabilize (resp., $\xrightarrow{\text{tr},k}$ -stabilize) to $\eta^k(\lambda)$. The following is our second main result, which we prove in §13 and §16.

Theorem 1.2.

- For any root system Φ , $L_\lambda^{\text{sym}}(k)$ is a polynomial in k with integer coefficients.
- For any simply laced root system Φ , $L_\lambda^{\text{tr}}(k)$ is a polynomial in k with integer coefficients.

We conjecture for all root systems Φ that these functions are polynomials in k with nonnegative integer coefficients. We call these polynomials *Ehrhart-like* because they count the number of points in some discrete region as it is dilated, but we note that in general the set of weights with given stabilization is not the set of lattice points of any convex polytope, or indeed any convex set (although these Ehrhart-like polynomials do include the usual Ehrhart polynomials of regular permutohedra).

That these Ehrhart-like polynomials apparently have nonnegative integer coefficients suggests that our interval-firing processes may have a deeper connection to the representation theory or algebraic geometry associated to the root system Φ , although we have no precise idea of what such a connection would be. There is some similarity between our interval-firing processes and the space of *quasi-invariants* of the Weyl group (see [ES03]). We thank Pavel Etingof for pointing this out to us.

As for possible connections to algebraic geometry: one can see in the above definitions of the interval-firing processes that rather than record the intervals corresponding to the values of $\langle \lambda, \alpha^\vee \rangle$ at which we allow firing, we recorded the intervals corresponding to the values of $\langle \lambda, \alpha^\vee \rangle + 1 = \langle \lambda + \frac{\alpha}{2}, \alpha^\vee \rangle$ at which we allow firing. This turns out to be more natural in many respects. And with this convention, the intervals defining the symmetric and truncated interval-firing processes are exactly the same as the intervals defining the (*extended*) Φ^\vee -Catalan and (*extended*) Φ^\vee -Shi hyperplane arrangements [PS00, Ath00]. The Catalan and Shi arrangements are known to have many remarkable combinatorial and algebraic properties, such as *freeness* [ER96, Ter02, Yos04]. Although we have no precise statement to this effect, empirically it seems that many of the remarkable properties of these families of hyperplane arrangements are reflected in the interval-firing processes. See Remark 11.8 for more discussion of connections with hyperplane arrangements.

Finally, we remark that a kind of “chip-firing for root systems” was recently studied by Benkart, Klivans, and Reiner [BKR18]. However, what Benkart-Klivans-Reiner studied was in fact M-matrix chip-firing with respect to the Cartan matrix \mathbf{C} of the root system Φ . As we discuss later (see §10), this Cartan matrix chip-firing is analogous to root-firing *where we only allow firing of the simple roots of Φ* . The root-firing processes we study in this paper allow firing of all the positive roots of Φ . Hence, our

set-up is quite different than the set-up of Benkart-Klivans-Reiner: for instance, the simple roots are always linearly independent, but there are many linear dependencies among the positive roots. Establishing confluence for Cartan matrix chip-firing is easy since the fact that the simple roots are pairwise non-acute implies confluence holds locally; whereas two positive roots may form an acute angle and hence confluence for interval-firing processes is a much more delicate question. Nevertheless, we do explain in Remark 10.3 how Cartan matrix chip-firing can be obtained from our interval-firing processes by taking a $k \rightarrow \infty$ limit.

Now let us outline the rest of the paper. In Part 1 we prove that the symmetric and truncated interval-firing processes are confluent. To do this, we first identify some Weyl group symmetries for both of the interval-firing processes (Theorem 5.1); in particular, we demonstrate that symmetric interval-firing is invariant under the action of the whole Weyl group (explaining its name). We then introduce the map η and explain how it labels the stable points for symmetric interval-firing (Lemma 6.6). We proceed to prove some polytopal results: we establish the aforementioned formula for traverse lengths of permutohedra (Theorem 7.6); this traverse length formula leads directly to a “permutohedron non-escaping lemma” (Lemma 8.2) which says that interval-firing processes get “trapped” inside of certain permutohedra. The confluence of symmetric interval-firing (Corollary 9.2) follows easily from the permutohedron non-escaping lemma. Finally, we establish the confluence of truncated interval-firing (Corollary 11.5) by first explaining how the map η also labels the stable points in the truncated case (Lemma 11.1), and then combining the permutohedron non-escaping lemma with a careful analysis of truncated interval-firing in rank 2.

In Part 2 we study the Ehrhart-like polynomials. We establish the existence of the symmetric Ehrhart-like polynomials (Theorem 13.3) via some basic Ehrhart theory for zonotopes (see, e.g., Theorem 13.1). Then, to establish the existence of the truncated Ehrhart-like polynomials in the simply laced case (Theorem 16.1), we study in detail the relationship between symmetric and truncated interval-firing and in particular how the connected components of the graphs of these processes “decompose” into smaller connected components in a way consistent with the labeling map η (see §15). In the final section, §17, we explain how these Ehrhart-like polynomials also count the sizes of fibers of iterates of a certain operator on the weight lattice, another surprising property of these polynomials that would be worth investigating further.

Acknowledgements: We thank Jim Propp, both for several useful conversations and because his introduction of labeled chip-firing and his infectious enthusiasm for exploring its properties launched this project. The second author was supported by NSF grant #1122374.

Part 1. Confluence of symmetric and truncated interval-firing

2. BACKGROUND ON ROOT SYSTEMS

Here we review the basic facts about root systems we will need in the study of certain vector-firing processes we define in terms of a fixed root system Φ . For details, consult [Hum72], [Bou02], or [BB05].

Fix V , an n -dimensional real vector space with inner product $\langle \cdot, \cdot \rangle$. For a nonzero vector $\alpha \in V \setminus \{0\}$ we define its *covector* to be $\alpha^\vee := \frac{2\alpha}{\langle \alpha, \alpha \rangle}$. Then we define the *reflection* across the hyperplane orthogonal to α to be the linear map $s_\alpha: V \rightarrow V$ given by $s_\alpha(v) := v - \langle v, \alpha^\vee \rangle \alpha$.

Definition 2.1. A *root system* is a finite collection $\Phi \subseteq V \setminus \{0\}$ of nonzero vectors such that:

- (1) $\text{Span}_{\mathbb{R}}(\Phi) = V$;
- (2) $s_\alpha(\Phi) = \Phi$ for all $\alpha \in \Phi$;
- (3) $\text{Span}_{\mathbb{R}}(\{\alpha\}) \cap \Phi = \{\pm\alpha\}$ for all $\alpha \in \Phi$;
- (4) $\langle \beta, \alpha^\vee \rangle \in \mathbb{Z}$ for all $\alpha, \beta \in \Phi$.

We remark that sometimes the third condition is omitted and those root systems satisfying the third condition are called *reduced*. On the other hand, sometimes the fourth condition is omitted and those root systems satisfying the fourth condition are called *crystallographic*. We will assume that all root systems under consideration are reduced and crystallographic and from now on will drop these adjectives.

From now on in the paper we will fix a root system Φ in V . The vectors $\alpha \in \Phi$ are called *roots*. The dimension of V (which is n) is called the *rank* of the root system. The vectors α^\vee for $\alpha \in \Phi$ are called *coroots* and the set of coroots forms another root system, denoted Φ^\vee , in V .

We use W to denote the *Weyl group* of Φ , which is the subgroup of $GL(V)$ generated by the reflections s_α for $\alpha \in \Phi$. By the first and second conditions of the definition of a root system, W is isomorphic as an abstract group to a subgroup of the symmetric group on Φ , and hence is finite. Observe that the Weyl group of Φ^\vee is equal to the Weyl group of Φ . Also note that all transformations in W are orthogonal.

It is well-known that we can choose a set $\Delta \subseteq \Phi$ of *simple roots* which form a basis of V , and which divide the root system $\Phi = \Phi^+ \cup \Phi^-$ into *positive* roots Φ^+ and *negative* roots $\Phi^- := -\Phi^+$ so that any positive root $\alpha \in \Phi^+$ is a nonnegative integer combination of simple roots. The choice of Δ is equivalent to the choice of Φ^+ ; one way to choose Φ^+ is to choose a generic linear form and let Φ^+ be the set of roots which are positive according to this form. There are many choices for Δ but they are all conjugate under W . From now on we will fix a set of simple roots Δ , and thus also a set of positive roots Φ^+ . It is known that any $\alpha \in \Phi$ appears in some choice of simple roots (in fact, every $\alpha \in \Phi$ is W -conjugate to a simple root appearing with nonzero coefficient in its expansion in terms of simple roots) and hence $W(\Delta) = \Phi$. We use $\Delta = \{\alpha_1, \dots, \alpha_n\}$ to denote the simple roots with an arbitrary but fixed order. The coroots α_i^\vee for $i = 1, \dots, n$ are called the *simple coroots* and they of course form a set of simple roots for Φ^\vee . We will always make this choice of simple roots for the dual root system, unless stated otherwise. With this choice of simple roots for the dual root system, we have $(\Phi^\vee)^+ = (\Phi^+)^{\vee}$.

We use $\mathbf{C} := (\langle \alpha_i, \alpha_j^\vee \rangle) \in \mathbb{Z}^{n \times n}$ to denote the *Cartan matrix* of Φ . Clearly one can recover the root system Φ from the Cartan matrix \mathbf{C} , which is encoded by its Dynkin diagram. The *Dynkin diagram* of Φ is the graph with vertex set $[n] := \{1, 2, \dots, n\}$ obtained as follows: first for all $1 \leq i < j \leq n$ we draw $\langle \alpha_i, \alpha_j^\vee \rangle \langle \alpha_j, \alpha_i^\vee \rangle$ edges between i

and j ; then, if $\langle \alpha_i, \alpha_j^\vee \rangle \langle \alpha_j, \alpha_i^\vee \rangle \notin \{0, 1\}$ for some i and j , we draw an arrow on top of the edges between them, from i to j if $|\alpha_i| > |\alpha_j|$. If there are no arrows in the Dynkin diagram of Φ then we say that Φ is *simply laced*.

There are two important lattices related to Φ , the *root lattice* $Q := \text{Span}_{\mathbb{Z}}(\Phi)$ and the *weight lattice* $P := \{v \in V : \langle v, \alpha^\vee \rangle \in \mathbb{Z} \text{ for all } \alpha \in \Phi\}$. The elements of P are called the *weights* of Φ . By the assumption that Φ is crystallographic, we have $Q \subseteq P$. We use $\Omega := \{\omega_1, \dots, \omega_n\}$ to denote the dual basis to the basis of simple coroots $\{\alpha_1^\vee, \dots, \alpha_n^\vee\}$ (in other words, the ω_i are defined by $\langle \omega_i, \alpha_j^\vee \rangle = \delta_{i,j}$); the elements of Ω are called *fundamental weights*. Observe that $Q = \text{Span}_{\mathbb{Z}}(\Delta)$ and $P = \text{Span}_{\mathbb{Z}}(\Omega)$.

We use $P_{\geq 0}^{\mathbb{R}} := \text{Span}_{\mathbb{R}_{\geq 0}}(\Omega)$, $P_{\geq 0} := \text{Span}_{\mathbb{Z}_{\geq 0}}(\Omega)$ and similarly $Q_{\geq 0}^{\mathbb{R}} := \text{Span}_{\mathbb{R}_{\geq 0}}(\Delta)$, $Q_{\geq 0} := \text{Span}_{\mathbb{Z}_{\geq 0}}(\Delta)$. Note that $P_{\geq 0}^{\mathbb{R}}$ and $Q_{\geq 0}^{\mathbb{R}}$ are dual cones; moreover, because the simple roots are pairwise non-acute, we have $P_{\geq 0}^{\mathbb{R}} \subseteq Q_{\geq 0}^{\mathbb{R}}$. The elements of $P_{\geq 0}$ are called *dominant weights*. For every $\lambda \in P$ there exists a unique element in $W(\lambda) \cap P_{\geq 0}$ and we use λ_{dom} to denote this element. A dominant weight of great importance is the *Weyl vector* $\rho := \sum_{i=1}^n \omega_i$. It is well-known (and easy to check) that $\rho = \frac{1}{2} \sum_{\alpha \in \Phi^+} \alpha$.

The connected components of $\{v \in V : \langle v, \alpha^\vee \rangle \neq 0 \text{ for all } \alpha \in \Phi\}$ are called the *chambers* of Φ . The *fundamental chamber* is $C_0 := \{v \in V : \langle v, \alpha^\vee \rangle > 0 \text{ for all } \alpha \in \Phi^+\}$. The Weyl group acts freely and transitively on the chambers and hence every chamber is equal to wC_0 for some unique $w \in W$. Observe that $P_{\geq 0}^{\mathbb{R}}$ is the closure of C_0 .

If $U \subseteq V$ is any subspace spanned by roots, then $\Phi \cap U$ is a root system in U , which we call a *sub-root system* of Φ . The root lattice of $\Phi \cap U$ is of course $\text{Span}_{\mathbb{Z}}(\Phi \cap U)$ while the weight lattice is $\pi_U(P)$. Moreover, $\Phi^+ \cap U$ is a set of positive roots for $\Phi \cap U$, although $\Delta \cap U$ may not be a set of simple roots for $\Phi \cap U$. We will always consider the positive roots of $\Phi \cap U$ to be $\Phi^+ \cap U$ unless explicitly stated otherwise. The case of *parabolic sub-root systems* (where in fact $\Delta \cap U$ is a set of simple roots for $\Phi \cap U$) is of special significance: for $I \subseteq [n]$ we set $\Phi_I := \Phi \cap \text{Span}_{\mathbb{R}}(\{\alpha_i : i \in I\})$.

If there exists an orthogonal decomposition $V = V_1 \oplus V_2$ with $\emptyset \subsetneq V_1, V_2 \subsetneq V$ such that $\Phi = \Phi_1 \cup \Phi_2$ with $\Phi_i \subseteq V_i$ for $i = 1, 2$, then we write $\Phi = \Phi_1 \times \Phi_2$ and we say the root system Φ is *reducible*. Otherwise we say that it is *irreducible*. (Let us also declare by fiat that the empty set, although it is a root system, is not irreducible.) In other words, a root system is irreducible if and only if its Dynkin diagram is connected. The famous *Cartan-Killing classification* classifies all irreducible root systems up to isomorphism, where an *isomorphism* of root systems is a bijection between roots induced from an invertible orthogonal map, potentially composed with a global rescaling of the inner product. Figure 2 shows the Dynkin diagrams of all the irreducible root systems: these are the classical infinite series A_n for $n \geq 1$, B_n for $n \geq 2$, C_n for $n \geq 3$, D_n for $n \geq 4$, together with the exceptional root systems G_2 , F_4 , E_6 , E_7 , and E_8 . Our numbering of the simple roots is consistent with Bourbaki [Bou02]. In every case the subscript in the name of the root system denotes the number of nodes of the Dynkin diagram, which is also the number of simple roots, that is, the rank of Φ . These labels A_n , B_n , etc. are the *type* of the root system; we may also talk about, e.g., ‘‘Type A’’ root systems.

All constructions that depend on the root system Φ decompose in a simple way as a direct product of irreducible factors. Hence without loss of generality we will from now on **assume that Φ is irreducible**.




FIGURE 2. Dynkin diagrams of all irreducible root systems. The nodes corresponding to minuscule weights are filled in.

In an irreducible root system, there are at most two values of lengths $|\alpha|$ among the roots $\alpha \in \Phi$. Those roots whose lengths achieve the maximum value are called *long*, and those which do not are called *short*. The Weyl group W acts transitively on the long roots, and it also acts transitively on the short roots.

There is a natural partial order on P called the *root order* whereby $\mu \leq \lambda$ for $\mu, \lambda \in P$ if $\lambda - \mu \in Q_{\geq 0}$. When restricted to Φ^+ , this partial order is graded by *height*; the height of $\alpha = \sum_{i=1}^n c_i \alpha_i \in \Phi$ is $\sum_{i=1}^n c_i$. Because we have assumed that Φ is irreducible, there is a unique maximal element of Φ^+ according to root order, denoted θ and called the *highest root*. The highest root is always long. We use $\hat{\theta}$ to denote the unique (positive) root such that $\hat{\theta}^\vee$ is the highest root of the dual root system Φ^\vee (with respect to the choice of $\{\alpha_1^\vee, \dots, \alpha_n^\vee\}$ as simple roots). If Φ is simply laced then $\theta = \hat{\theta}$ and θ is the unique root which is a dominant weight; if Φ is not simply laced then θ and $\hat{\theta}$ are the two roots which are dominant weights. In the non-simply laced case we call $\hat{\theta}$ the *highest short root*: it is the maximal short root with respect to the root ordering.

The root lattice Q is a full rank sublattice of P ; hence the quotient P/Q is some finite abelian group. Note that $P/Q \simeq \text{coker}(\mathbf{C}^t)$ where we view the transposed matrix as a map $\mathbf{C}^t: \mathbb{Z}^n \rightarrow \mathbb{Z}^n$. The order of this group is called the *index of connection* of Φ and is denoted $f := |P/Q|$. There is a nice choice of coset representatives of P/Q , which we now describe. A dominant, nonzero weight $\lambda \in P_{\geq 0} \setminus \{0\}$ is called *minuscule* if $\langle \lambda, \alpha^\vee \rangle \in \{-1, 0, 1\}$ for all $\alpha \in \Phi$. Let us use Ω_m to denote the set of minuscule weights. Note that $\Omega_m \subseteq \Omega$, i.e., a minuscule weight must be a fundamental weight. In Figure 2, the vertices corresponding to minuscule weights are filled in. In fact, there are $f - 1$ minuscule weights and the minuscule weights together with zero form a collection

of coset representatives of P/Q . We use $\Omega_m^0 := \Omega_m \cup \{0\}$ to denote the set of these representatives.

There is another characterization of minuscule weights that we will find useful. Namely, for a dominant weight $\lambda \in P_{\geq 0}$ we have that $\lambda \in \Omega_m^0$ if and only if λ is the minimal element according to root order in $(Q + \lambda) \cap P_{\geq 0}$.

This last characterization of minuscule weight can also be described in terms of certain polytopes called *(W)-permutohedra*. Permutohedra will play a key role for us in our understanding of interval-firing processes, so let us review these now. For $v \in V$, we define the *permutohedron* associated to v to be $\Pi(v) := \text{ConvexHull } W(v)$, a convex polytope in V . And for a weight $\lambda \in P$, we define $\Pi^Q(\lambda) := \Pi(\lambda) \cap (Q + \lambda)$, which we call the *discrete permutohedron* associated to λ .

The following simple proposition describes the containment of permutohedra (see also [Ste98, 1.2]):

Proposition 2.2. *For $u, v \in P_{\geq 0}^{\mathbb{R}}$ we have $\Pi(u) \subseteq \Pi(v)$ if and only if $v - u \in Q_{\geq 0}^{\mathbb{R}}$. Hence for $\mu, \lambda \in P_{\geq 0}$ we have $\Pi^Q(\mu) \subseteq \Pi^Q(\lambda)$ if and only if $\mu \leq \lambda$ (in root order).*

Proof. First suppose that u and v are strictly inside the fundamental chamber C_0 , i.e., that we have $\langle u, \alpha_i^\vee \rangle > 0$ and $\langle v, \alpha_i^\vee \rangle > 0$ for all $i \in [n]$. By the *inner cone* of polytope at a vertex, we mean the affine convex cone spanned by the edges of the polytope incident to that vertex in the direction “outward” from that vertex. Note that a point belongs to a polytope if and only if it belongs to the inner cone of that polytope at every vertex. Since the walls of the fundamental chamber are orthogonal to the simple roots, it is easy to see that if u and v are strictly inside the fundamental chamber then the inner cone of $\Pi(u)$ at u is spanned by the negatives of the simple roots, and ditto for the inner cone of $\Pi(v)$ and v . So if we do not have $v - u \in Q_{\geq 0}^{\mathbb{R}}$, then clearly u does not belong to $\Pi(v)$. Hence suppose that $v - u \in Q_{\geq 0}^{\mathbb{R}}$. Every vertex of $\Pi(u)$ belongs to the inner cone of $\Pi(u)$ at u ; i.e., $u - u' \in Q_{\geq 0}^{\mathbb{R}}$ for all $u' \in W(u)$. Thus for all $u' \in W(u)$ we have $v - u' \in Q_{\geq 0}^{\mathbb{R}}$; i.e., every point in $\Pi(u)$ is in the inner cone of $\Pi(v)$ at v . But then by the W -invariance of permutohedra, we conclude that every point in $\Pi(u)$ is in the inner cone of $\Pi(v)$ at every vertex of $\Pi(v)$, and hence that $\Pi(u) \subseteq \Pi(v)$, as claimed.

For arbitrary $u, v \in P_{\geq 0}^{\mathbb{R}}$, note $\Pi(u) = \bigcap_{\varepsilon > 0} \Pi(u + \varepsilon\rho)$ and $\Pi(v) = \bigcap_{\varepsilon > 0} \Pi(v + \varepsilon\rho)$, and $u + \varepsilon\rho$ and $v + \varepsilon\rho$ will be strictly inside the fundamental chamber for all $\varepsilon > 0$. Thus the result for arbitrary $u, v \in P_{\geq 0}^{\mathbb{R}}$ follows from the preceding paragraph. \square

So in light of Proposition 2.2, we see that minuscule weights can also be characterized as follows: for $\lambda \in P_{\geq 0}$ we have $\lambda \in \Omega_m^0$ if and only if $\Pi^Q(\lambda) = W(\lambda)$. For references for all these various characterizations of and facts about minuscule weights, see [BKR18, Proposition 3.10] (who in particular credit Stembridge [Ste98] for some of these facts).

3. BACKGROUND ON BINARY RELATIONS AND CONFLUENCE

Interval-firing will formally be defined to be a binary relation on the weight lattice of Φ . Before giving the precise definition, we review some general notation and results concerning binary relations. Let X be a set and \rightarrow a binary relation on X . We use $\Gamma \rightarrow$ to denote the directed graph (from now on, “digraph”) with vertex set X and

with a directed edge (x, y) whenever $x \rightarrow y$. Clearly Γ_{\rightarrow} contains exactly the same information as \rightarrow and we will often implicitly identify binary relations and digraphs (specifically, digraphs without multiple edges in the same direction) in this way. We use $\xrightarrow{*}$ to denote the reflexive, transitive closure of \rightarrow : that is, we write $x \xrightarrow{*} y$ to mean that $x = x_0 \rightarrow x_1 \rightarrow \cdots \rightarrow x_k = y$ for some $k \in \mathbb{Z}_{\geq 0}$. In other words, $x \xrightarrow{*} y$ means there is a path from x to y in Γ_{\rightarrow} . We use \leftrightarrow to denote the symmetric closure of \rightarrow : $x \leftrightarrow y$ means that $x \rightarrow y$ or $y \rightarrow x$. For any digraph Γ , we use Γ^{un} to denote the underlying undirected graph of Γ ; in fact, we view Γ^{un} as a digraph: it has edges (x, y) and (y, x) whenever (x, y) is an edge of Γ . Hence $\Gamma_{\leftrightarrow} = \Gamma^{\text{un}}$. Finally, we use $\xleftrightarrow{*}$ to denote the reflexive, transitive, symmetric closure of \rightarrow : $x \xleftrightarrow{*} y$ means that $x = x_0 \leftrightarrow x_1 \leftrightarrow \cdots \leftrightarrow x_k = y$ for some $k \in \mathbb{Z}_{\geq 0}$. In other words, $x \xleftrightarrow{*} y$ means there is a path from x to y in Γ^{un} .

Now let us review some notions of confluence for binary relations. Here we generally follow standard terminology in the theory of abstract rewriting systems, as laid out for instance in [Hue80]; however, instead following chip-firing terminology, we use “stable” in place of what would normally be called “irreducible,” and rather than “normal forms” we refer to “stabilizations.” We say that \rightarrow is *terminating* (also sometimes called *noetherian*) if there is no infinite sequence of relations $x_0 \rightarrow x_1 \rightarrow x_2 \rightarrow \cdots$; i.e., \rightarrow is terminating means that Γ_{\rightarrow} has no infinite paths (which implies in particular that this digraph has no directed cycles). Generally speaking, the relations we are most interested in will all be terminating and it will be easy for us to establish that they are terminating. For $x \in X$, we say that \rightarrow is *confluent from x* if whenever $x \xrightarrow{*} y_1$ and $x \xrightarrow{*} y_2$, there is y_3 such that $y_1 \xrightarrow{*} y_3$ and $y_2 \xrightarrow{*} y_3$. We say $x \in X$ is \rightarrow -*stable* (or just *stable* if the context is clear) if there is no $y \in X$ with $x \rightarrow y$. In graph-theoretic language, x is \rightarrow -stable means that x is a sink (vertex of outdegree zero) of Γ_{\rightarrow} . If \rightarrow is terminating, then for every $x \in X$ there must be at least one stable $y \in X$ with $x \xrightarrow{*} y$. On the other hand, if \rightarrow is confluent from $x \in X$, then there can be at most one stable $y \in X$ with $x \xrightarrow{*} y$. Hence if \rightarrow is terminating and is confluent from x , then there exists a unique stable y with $x \xrightarrow{*} y$; we call this y the \rightarrow -*stabilization* (or just *stabilization* if the context is clear) of x . We say that \rightarrow is *confluent* if it is confluent from every $x \in X$. As we just explained, if \rightarrow is confluent and terminating then a unique stabilization of x exists for all $x \in X$. A weaker notion than confluence is that of local confluence: we say that \rightarrow is *locally confluent* if for any $x \in X$, if $x \rightarrow y_1$ and $x \rightarrow y_2$, then there is some y_3 with $y_1 \xrightarrow{*} y_3$ and $y_2 \xrightarrow{*} y_3$. Figure 3 gives some examples of relations comparing these various notions of confluence and termination. Observe that there is no example in this figure of a relation that is locally confluent and terminating but not confluent. That is no coincidence: Newman’s lemma, a.k.a. the diamond lemma, says that local confluence plus termination implies confluence.

Lemma 3.1 (Diamond lemma, see [New42, Theorem 3] or [Hue80, Lemma 2.4]). *Suppose \rightarrow is terminating. Then \rightarrow is confluent if and only if it is locally confluent.*

4. DEFINITION OF INTERVAL-FIRING

In this section we formally define the interval-firing processes in their most general form. We use the notation $\mathbf{k} \in \mathbb{Z}[\Phi]^W$ to mean that \mathbf{k} is an integer-valued function on




FIGURE 3. Examples of various relations: (I) is confluent from x but not from y ; (II) and (III) are confluent but not terminating; (IV) and (VI) are locally confluent but not confluent; (V) is confluent and terminating.

the roots of Φ that is invariant under the action of the Weyl group. We write $\mathbf{a} \leq \mathbf{b}$ to mean that $\mathbf{a}(\alpha) \leq \mathbf{b}(\alpha)$ for all $\alpha \in \Phi$. We use the notation $\mathbf{k} = k$ to mean that \mathbf{k} is constantly equal to k . We also use the obvious notation $a\mathbf{a} + b\mathbf{b}$ for linear combinations of these functions. We use $\mathbb{N}[\Phi]^W$ to denote the set of $\mathbf{k} \in \mathbb{Z}[\Phi]^W$ with $\mathbf{k} \geq 0$. We write $\rho_{\mathbf{k}} := \sum_{i=1}^n \mathbf{k}(\alpha_i)\omega_i$. Since we have assumed that Φ is irreducible, there are at most two W -orbits of Φ : the short roots and the long roots. If Φ is simply laced then it has a single Weyl group orbit and $\mathbf{k} = k$ for some constant $k \in \mathbb{Z}$; otherwise, we have two constants $k_s, k_l \in \mathbb{Z}$ so that $\mathbf{k}(\alpha) = k_s$ if α is short and $\mathbf{k}(\alpha) = k_l$ if α is long.

For $\mathbf{k} \in \mathbb{N}[\Phi]^W$, the *symmetric interval-firing process* is the binary relation $\xrightarrow{\text{sym}, \mathbf{k}}$ on P defined by

$$\lambda \xrightarrow{\text{sym}, \mathbf{k}} \lambda + \alpha, \text{ for } \lambda \in P \text{ and } \alpha \in \Phi^+ \text{ with } \langle \lambda + \frac{\alpha}{2}, \alpha^\vee \rangle \in [-\mathbf{k}(\alpha), \mathbf{k}(\alpha)]$$

and the *truncated interval-firing process* is the binary relation $\xrightarrow{\text{tr}, \mathbf{k}}$ on P defined by

$$\lambda \xrightarrow{\text{tr}, \mathbf{k}} \lambda + \alpha, \text{ for } \lambda \in P \text{ and } \alpha \in \Phi^+ \text{ with } \langle \lambda + \frac{\alpha}{2}, \alpha^\vee \rangle \in [-\mathbf{k}(\alpha) + 1, \mathbf{k}(\alpha)].$$

From now on we will often think about a relation \rightarrow as $\Gamma \rightarrow$. So we use the shorthand notations $\Gamma_{\text{sym}, \mathbf{k}} := \Gamma \xrightarrow{\text{sym}, \mathbf{k}}$ and $\Gamma_{\text{tr}, \mathbf{k}} := \Gamma \xrightarrow{\text{tr}, \mathbf{k}}$.

Example 4.1. The irreducible rank 2 root systems are A_2 , B_2 and G_2 . The positive roots and fundamental weights for these root systems are depicted in Figure 4. In Figures 5, 6, and 7 we depict the the truncated and symmetric interval-firing processes $\Gamma_{\text{tr}, \mathbf{k}}$ and $\Gamma_{\text{sym}, \mathbf{k}}$ for $\mathbf{k} = 0, 1, 2$ for these three root systems. Of course these graphs are infinite, so we depict the “interesting part” of the graphs near the origin (which is circled in black). The colors in these drawings correspond to classes of weights modulo the root lattice (hence there are three colors in the A_2 graphs, two in the B_2 graphs, and one in the G_2 graphs). Note that as \mathbf{k} increases, the scale of the drawing is






FIGURE 4. The positive roots of the rank 2 root systems A_2 , B_2 , and G_2 . The elements of $\Omega \cup \{0\}$ are shown in red.

not maintained. Most, if not all, of the features of truncated and symmetric interval-firing that we care about are visible already in rank 2. Thus the reader is encouraged, while reading the rest of this paper, to return to these figures and understand how each of the results apply to these two dimensional examples.

Remark 4.2. Let us recall Propp’s labeled chip-firing process (studied in [HMP17]), which motivated our study of interval-firing processes. The states of labeled chip-firing are configurations of labeled chips on the infinite path graph \mathbb{Z} , such as:



If two chips with labels occupy the same position, we may *fire* them, which sends the lesser-labeled chip one vertex to the right and the greater-labeled chip one vertex to the left. For instance, firing the chips $\textcircled{1}$ and $\textcircled{2}$ above leads to



Firing the chips \textcircled{i} and \textcircled{j} with $i < j$ corresponds to $c \rightarrow c + (e_i - e_j)$, where the integer vector $c := (c_1, \dots, c_N) \in \mathbb{Z}^N$ is given by $c_i :=$ the position of the chip \textcircled{i} . In this way central-firing (the subject of our sequel paper [GHMP17]) is the same as the labeled chip-firing process for Φ of Type A. Via this same correspondence between lattice vectors and configurations of chips, symmetric and truncated interval-firing in Type A can also be seen as “labeled chip-firing processes” that consist of the same chip-firing moves, which send chip \textcircled{i} one vertex to the right and chip \textcircled{j} one vertex to the left for any $i < j$, but where we allow these moves to be applied under different conditions: namely, when the position of chip \textcircled{i} minus the position of chip \textcircled{j} is either in the interval $[-k-1, k-1]$ (in the symmetric case) or in the interval $[-k, k-1]$ (in the




FIGURE 5. The graphs $\Gamma_{\text{tr},k}$ and $\Gamma_{\text{sym},k}$ for $\Phi = A_2$ and $k = 0, 1, 2$.

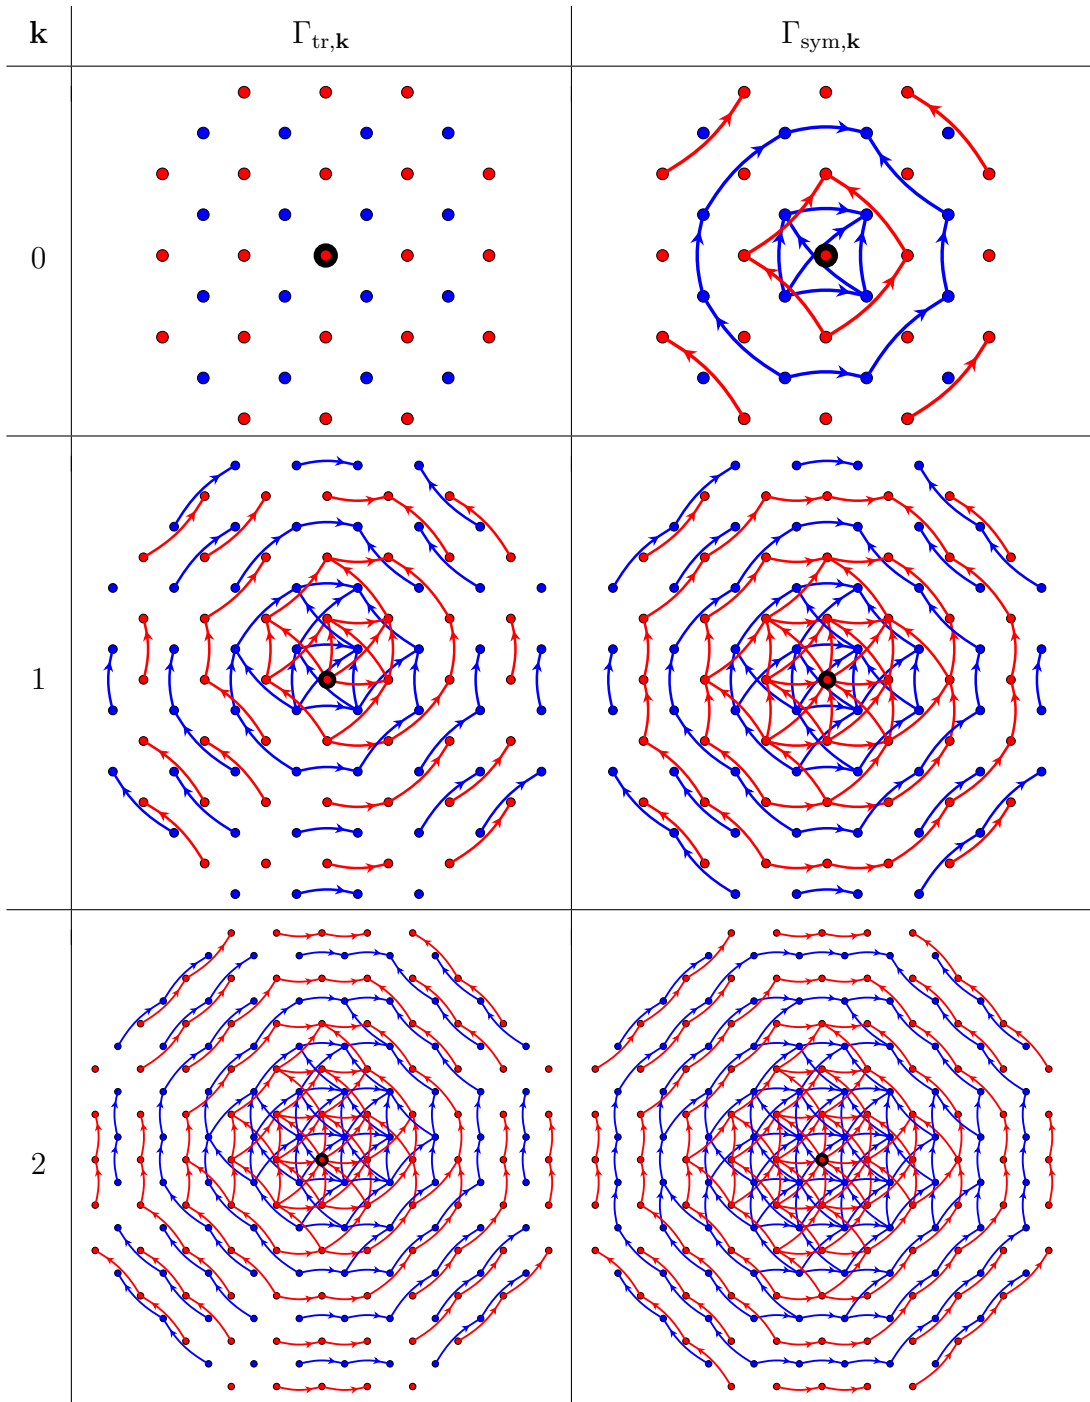


FIGURE 6. The graphs $\Gamma_{\text{tr},k}$ and $\Gamma_{\text{sym},k}$ for $\Phi = B_2$ and $k = 0, 1, 2$.

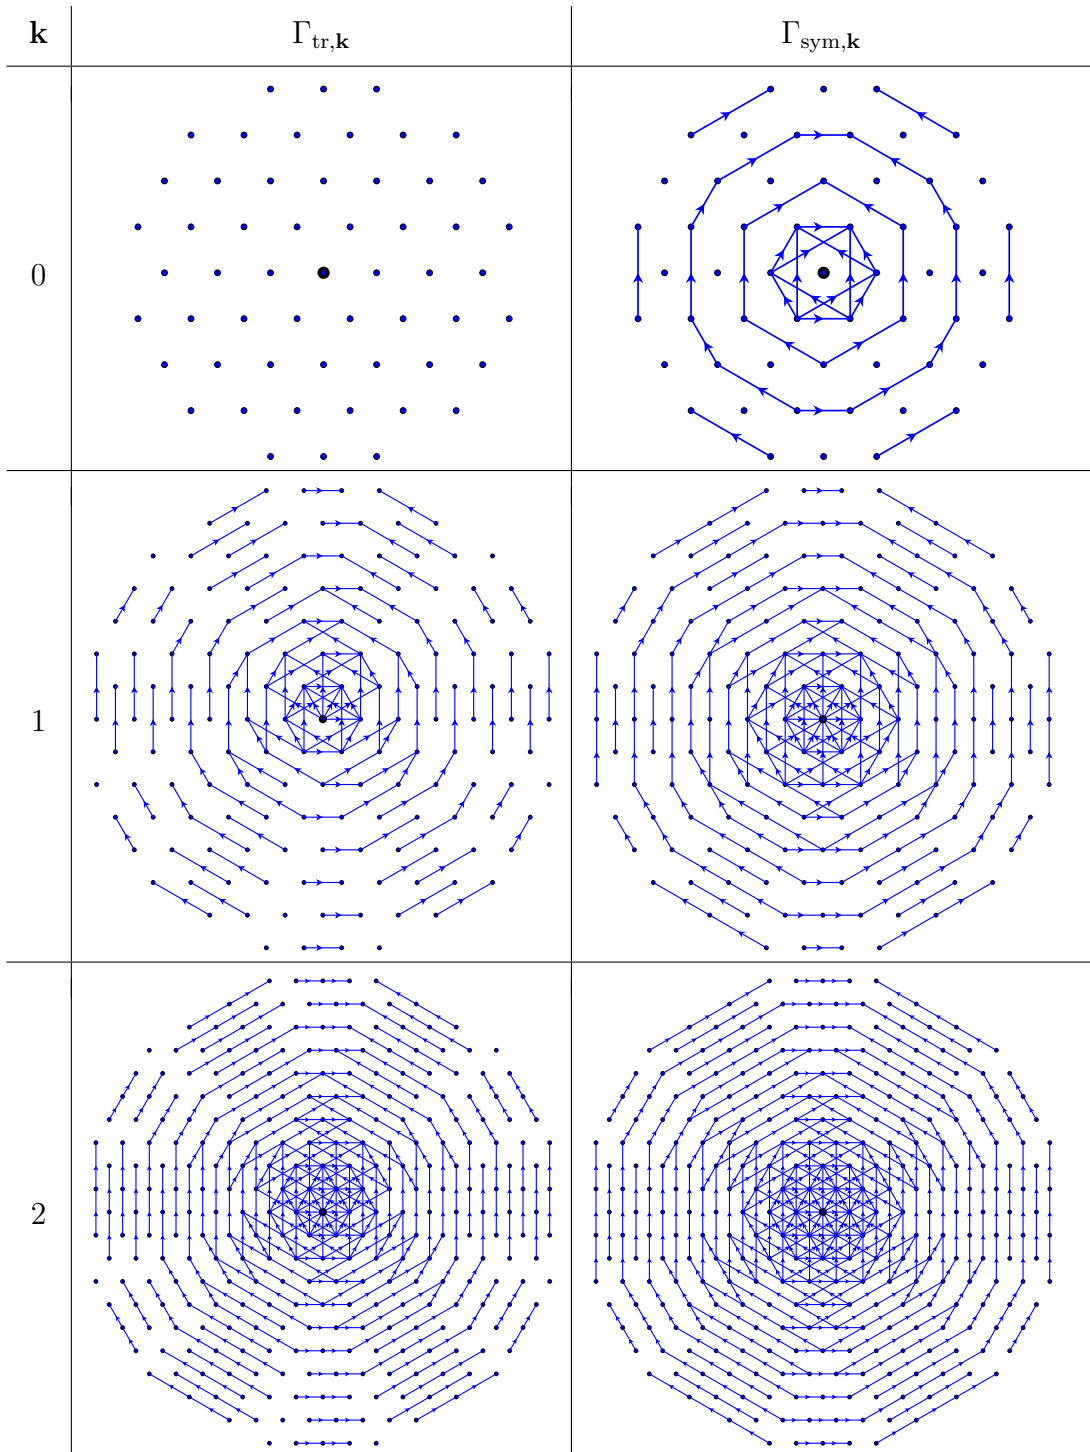




FIGURE 7. The graphs $\Gamma_{\text{tr},k}$ and $\Gamma_{\text{sym},k}$ for $\Phi = G_2$ and $k = 0, 1, 2$.


truncated case). For example, consider the smallest non-trivial case of these interval-firing processes, which is symmetric interval-firing with $k = 0$. This corresponds to the labeled chip-firing process that allows the transposition of the chips (i) and (j) with $i < j$ when (i) is one position to the left of (j) . It is immediately apparent that this process is confluent; for instance, the configuration




$\xrightarrow{\text{sym},0}$ -stabilizes to



In general the stabilization will weakly sort each collection of contiguous chips, while leaving the underlying unlabeled configuration of chips the same. The next smallest case to consider is truncated interval-firing with $k = 1$. This corresponds to the labeled chip-firing process that allows both the transposition moves from the symmetric $k = 0$ case, and the usual labeled chip-firing moves from the central-firing case. The reader can verify that for instance the configuration



$\xrightarrow{\text{tr},1}$ -stabilizes to



Here it is less obvious that confluence holds (although it is not too hard to prove this fact directly via a diamond lemma argument). The reader is now encouraged to experiment with this labeled chip-firing interpretation of symmetric and truncated interval-firing for higher values of k . Note that increasing k allows for the firing of chips (i) and (j) when they are further apart.

In our further treatment of the interval-firing processes we will focus on the geometric picture (on display in Example 4.1) and not the chip-firing picture (discussed in Remark 4.2).

To close out this section, let us demonstrate that the interval-firing processes are always terminating. This is straightforward because the collection Φ^+ of vectors we are adding is acyclic.

Proposition 4.3. *For $\mathbf{k} \in \mathbb{N}[\Phi]^W$, the relations $\xrightarrow{\text{sym},\mathbf{k}}$ and $\xrightarrow{\text{tr},\mathbf{k}}$ are terminating.*

Proof. It is enough to show this for $\xrightarrow{\text{sym},\mathbf{k}}$, which has more firing moves than $\xrightarrow{\text{tr},\mathbf{k}}$.

For $\lambda \in P$ define $\varphi(\lambda) := \langle \rho_{\mathbf{k}+1} - \lambda, \rho_{\mathbf{k}+1} - \lambda \rangle$; in other words, $\varphi(\lambda)$ is the length of the

vector $\rho_{\mathbf{k}+1} - \lambda$. Suppose $\lambda \xrightarrow{\text{sym}, \mathbf{k}} \lambda + \alpha$ for $\alpha \in \Phi^+$. Then,

$$\begin{aligned} \varphi(\lambda) - \varphi(\lambda + \alpha) &= \langle \rho_{\mathbf{k}+1} - \lambda, \rho_{\mathbf{k}+1} - \lambda \rangle - \langle \rho_{\mathbf{k}+1} - (\lambda + \alpha), \rho_{\mathbf{k}+1} - (\lambda + \alpha) \rangle \\ &= 2\langle \rho_{\mathbf{k}+1}, \alpha \rangle - 2\langle \lambda, \alpha \rangle \\ &\geq \frac{4}{\langle \alpha, \alpha \rangle} (\mathbf{k}(\alpha) + 1 - \mathbf{k}(\alpha)) = \frac{4}{\langle \alpha, \alpha \rangle}, \end{aligned}$$

where we use the fact that $\langle \lambda, \alpha \rangle \leq \frac{2}{\langle \alpha, \alpha \rangle} \mathbf{k}(\alpha)$ since $\lambda \xrightarrow{\text{sym}, \mathbf{k}} \lambda + \alpha$. So each firing move causes the quantity $\varphi(\lambda)$ to decrease by at least some fixed nonzero amount. But we have $\varphi(\lambda) \geq 0$ because it is measuring the length of a vector. Thus indeed $\xrightarrow{\text{sym}, \mathbf{k}}$ is terminating. \square

5. SYMMETRIES OF INTERVAL-FIRING PROCESSES

In this section we study the symmetries of the two interval-firing processes. Since the set of positive roots Φ^+ is an “oriented” set of vectors, we do not expect the directed graphs $\Gamma_{\text{sym}, \mathbf{k}}$ and $\Gamma_{\text{tr}, \mathbf{k}}$ to have many symmetries, and certainly none coming from the Weyl group. But if we consider instead the undirected graphs $\Gamma_{\text{sym}, \mathbf{k}}^{\text{un}}$ and $\Gamma_{\text{tr}, \mathbf{k}}^{\text{un}}$ (corresponding to the symmetric relations $\longleftrightarrow_{\text{sym}, \mathbf{k}}$ and $\longleftrightarrow_{\text{tr}, \mathbf{k}}$), we will see that both of these do in fact have symmetries coming from the Weyl group.

For the symmetric interval-firing process, the graph $\Gamma_{\text{sym}, \mathbf{k}}^{\text{un}}$ is invariant under the action of the whole Weyl group W . This explains the name “symmetric” for the process: it has the biggest possible group of symmetries. As for the truncated process, in order to understand its symmetries we need to introduce a certain subgroup of the Weyl group $C \subseteq W$. In fact this C is an abelian group and satisfies $C \simeq P/Q$. In our definition of C we follow Lam and Postnikov [LP12]². The *Coxeter number* of Φ , another fundamental invariant of the root system, is $h := \langle \rho, \widehat{\theta}^\vee \rangle + 1$. (The Coxeter number is also equal to $h = 1 + \sum_{i=1}^n a_i$ where $\theta = \sum_{i=1}^n a_i \alpha_i$.) Lam and Postnikov [LP12, §5] defined the subgroup $C := \{w \in W : \rho - w(\rho) \in hP\}$ of the Weyl group and explained (using the affine Weyl group, which we will not discuss here) that C is naturally isomorphic to P/Q : the isomorphism is explicitly given by $w \mapsto \omega \in \Omega_m^0$ if and only if $\rho - w(\rho) = h\omega$. (Since $\rho - w(\rho) \in Q$ for any $w \in W$, a consequence of this description of the isomorphism is that $h \cdot (P/Q) = \{0\}$.) As they mention, this subgroup was also studied before by Verma [Ver75], but in spite of its significance it does not seem to have any name other than C in the root system literature. Lam and Postnikov gave another characterization [LP12, Proposition 6.4] of C that will be useful for us: $C = \{w \in W : w(\{\alpha_0^\vee, \alpha_1^\vee, \dots, \alpha_n^\vee\}) = \{\alpha_0^\vee, \alpha_1^\vee, \dots, \alpha_n^\vee\}\}$, where we use the suggestive notation $\alpha_0^\vee := -\widehat{\theta}^\vee$.

Theorem 5.1. *Let $\mathbf{k} \in \mathbb{N}[\Phi]^W$. Set $\Gamma := \Gamma_{\text{sym}, \mathbf{k}}^{\text{un}}$ or $\Gamma := \Gamma_{\text{tr}, \mathbf{k}}^{\text{un}}$. Then,*

- *if $\Gamma = \Gamma_{\text{sym}, \mathbf{k}}^{\text{un}}$, the linear map $v \mapsto w(v)$ is an automorphism of Γ for all $w \in W$;*

²Lam and Postnikov worked in a completely dual setting to ours: that is, they described a copy of the coweight lattice modulo the coroot lattice inside of W ; hence, they used θ instead of $\widehat{\theta}$, etc.

- if $\Gamma = \Gamma_{\text{tr}, \mathbf{k}}^{\text{un}}$, the affine map $v \mapsto w(v - \frac{1}{h}\rho) + \frac{1}{h}\rho$ is an automorphism of Γ for all $w \in C \subseteq W$.

Proof. If $\Gamma = \Gamma_{\text{sym}, \mathbf{k}}^{\text{un}}$ set $c := 0$, and if $\Gamma = \Gamma_{\text{tr}, \mathbf{k}}^{\text{un}}$ set $c := 1$. Consider the hyperplane arrangement $\mathcal{H} := \left\{ H_{\alpha^\vee, \frac{c}{2}} : \alpha \in \Phi^+ \right\}$ with hyperplanes $H_{\alpha^\vee, \frac{c}{2}} := \{v \in V : \langle v, \alpha^\vee \rangle = \frac{c}{2}\}$.

First we claim that if for $w \in W$ and $u \in V$ the affine map $\varphi : v \mapsto w(v - u) + u$ is an automorphism of \mathcal{H} which maps P to P , then it is an automorphism of Γ (by an automorphism of the hyperplane arrangement, we mean an invertible affine map φ such that φ permutes the hyperplanes in \mathcal{H}). Indeed, observe that there is an edge in Γ between λ and μ if and only if there is some $\alpha \in \Phi^+$ such that $\mu = \lambda + \alpha$ and $\max(\{|\langle \mu, \alpha^\vee \rangle - \frac{c}{2}|, |\langle \lambda, \alpha^\vee \rangle - \frac{c}{2}|\}) \leq \mathbf{k}(\alpha) + 1 - \frac{c}{2}$. So suppose there is an edge between λ and μ in the α direction. Any φ of this form will satisfy $\varphi(\mu) - \varphi(\lambda) = w(\alpha)$ and $\varphi(H_{\alpha^\vee, \frac{c}{2}}) = H_{\pm w(\alpha)^\vee, \frac{c}{2}}$ (where the sign \pm is chosen so that $\pm w(\alpha) \in \Phi^+$). Moreover, since all Weyl group elements are orthogonal, and, in particular, preserve distances, the distance from μ to $H_{\alpha^\vee, \frac{c}{2}}$ will be the same as the distance from $\varphi(\mu)$ to $H_{\pm w(\alpha)^\vee, \frac{c}{2}}$, and ditto for λ . But $|\langle \mu, \alpha^\vee \rangle - \frac{c}{2}|$ is precisely the distance from μ to $H_{\alpha^\vee, \frac{c}{2}}$, and ditto for λ . Hence indeed we will get that $\varphi(\mu) = \varphi(\lambda) + w(\alpha)$ and that

$$\begin{aligned} \max\left(\left|\left\langle \varphi(\mu), (\pm w(\alpha))^\vee \right\rangle - \frac{c}{2}\right|, \left|\left\langle \varphi(\lambda), (\pm w(\alpha))^\vee \right\rangle - \frac{c}{2}\right|\right) \\ \leq \mathbf{k}(\alpha) + 1 - \frac{c}{2} = \mathbf{k}(\pm w(\alpha)) + 1 - \frac{c}{2}, \end{aligned}$$

which means there is an edge in Γ between $\varphi(\lambda)$ and $\varphi(\mu)$ in the $\pm w(\alpha)$ direction. To see that conversely if there is an edge between $\varphi(\lambda)$ and $\varphi(\mu)$ in Γ , there is one between λ and μ , use that φ is invertible and φ^{-1} is of the same form.

In the case $c = 0$, the hyperplane arrangement \mathcal{H} is just the Coxeter arrangement of Φ and it is easy to see that every $w \in W$ is an automorphism of \mathcal{H} .

Now consider the case $c = 1$, in which case \mathcal{H} is (a scaled version of) the Φ^\vee -Linial arrangement; see for instance [PS00] and [Ath00]. We claim that $\varphi : v \mapsto w(v - \frac{c}{h}\rho) + \frac{c}{h}\rho$ is an automorphism of \mathcal{H} for all $w \in C$. So suppose $x \in H_{\alpha^\vee, \frac{c}{2}}$; we want to show that $\varphi(x) \in H_{\pm w(\alpha)^\vee, \frac{c}{2}}$ where the sign \pm is chosen so that $\pm w(\alpha)$ is positive. (The reverse implication will then follow from consideration of $\varphi^{-1} = w^{-1}(v - \frac{c}{h}\rho) + \frac{c}{h}\rho$.) We have

$$(5.1) \quad \langle \varphi(x), w(\alpha)^\vee \rangle = \frac{c}{2} - \left\langle \frac{c}{h}\rho, \alpha^\vee \right\rangle + \left\langle \frac{c}{h}\rho, w(\alpha)^\vee \right\rangle.$$

Write $\alpha^\vee = \sum_{i=1}^n a_i \alpha_i^\vee$, with the convention $a_0 := 0$. By a result of Lam-Postnikov mentioned above, there is a permutation $\pi : \{0, 1, \dots, n\} \rightarrow \{0, 1, \dots, n\}$ such that $w(\alpha_i^\vee) = \alpha_{\pi(i)}^\vee$ (with the aforementioned convention $\alpha_0^\vee := -\hat{\theta}^\vee$ where $\hat{\theta}^\vee$ is the highest root of Φ^\vee). Thus, $w(\alpha)^\vee = \sum_{i=1}^n a_i \alpha_{\pi(i)}^\vee$.

We will consider two cases. First suppose that $a_{\pi^{-1}(0)} = 0$. Then $w(\alpha)^\vee$ is clearly a positive root, so $\pm = +$; moreover, we have $\langle \frac{c}{h}\rho, \alpha^\vee \rangle = \langle \frac{c}{h}\rho, w(\alpha)^\vee \rangle = \frac{c}{h} \cdot \sum_{i=1}^n a_i$. So from (5.1) we get that $\langle \varphi(x), w(\alpha)^\vee \rangle = \frac{c}{2}$, that is, $\varphi(x) \in H_{\pm w(\alpha)^\vee, \frac{c}{2}}$, as desired.

Now suppose that $a_{\pi^{-1}(0)} \neq 0$. We claim that this forces $a_{\pi^{-1}(0)} = 1$: indeed, otherwise the height of $w(\alpha)^\vee$ would be strictly less than $-(h-1)$, which is impossible

because $-\widehat{\theta}^\vee$ has height $-(h-1)$ and is the root in Φ^\vee of smallest height. So indeed we have $a_{\pi^{-1}(0)} = 1$. Note also that in this case the height of $w(\alpha)^\vee$ is a negative root and hence $w(\alpha)^\vee$ is negative, so $\pm = -$. Then we compute

$$-\left\langle \frac{c}{h}\rho, \alpha^\vee \right\rangle + \left\langle \frac{c}{h}\rho, w(\alpha)^\vee \right\rangle = -\left\langle \frac{c}{h}\rho, \alpha_{\pi^{-1}(0)}^\vee \right\rangle + \left\langle \frac{c}{h}\rho, \alpha_0^\vee \right\rangle = -\frac{c}{h} - \left(\frac{c}{h}(h-1) \right) = -c.$$

Thus from (5.1) we get that $\langle \varphi(x), -w(\alpha)^\vee \rangle = -\frac{c}{2} + c = \frac{c}{2}$, that is, $\varphi(x) \in H_{\pm w(\alpha)^\vee, \frac{c}{2}}$, as desired.

Finally, the description of C given above says that $\varphi(0) = w(0 - \frac{c}{h}\rho) + \frac{c}{h}\rho = c\omega$ for some $\omega \in \Omega_m^0$. Hence indeed φ maps P to P . \square

6. SINKS OF SYMMETRIC INTERVAL-FIRING AND THE MAP η

Recall that our overall strategy for proving confluence of the interval-firing processes is to show that they get “trapped” inside certain permutohedra, and then to analyze where these processes must terminate. In order to carry out this strategy, we need to understand what are the possible final points we terminate at, i.e., what are the stable points of these processes.

In this section we describe the $\xrightarrow{\text{sym}, \mathbf{k}}$ -stable points, i.e., the sinks of $\Gamma_{\text{sym}, \mathbf{k}}$. We will show in particular that there is a way to consistently label the sinks of $\Gamma_{\text{sym}, \mathbf{k}}$ across all values of \mathbf{k} .


In order to define this labeling we need to review some basic facts about parabolic subgroups and parabolic cosets. Recall that the Weyl group W is generated by the *simple reflections* $s_i := s_{\alpha_i}$ for $i = 1, \dots, n$. For any $w \in W$ we use $\ell(w)$ to denote the *length* of w , which is the length of the shortest representation of w as a product of simple reflections. An *inversion* of w is a positive root $\alpha \in \Phi^+$ for which $w(\alpha)$ is negative. The length $\ell(w)$ is equal to the number of inversions of w . The identity is the only Weyl group element of length zero. The simple reflections are the only Weyl group elements of length one: s_i sends α_i to $-\alpha_i$ and permutes $\Phi^+ \setminus \{\alpha_i\}$. A (*right*) *descent* of $w \in W$ is a simple reflection s_i such that $\ell(ws_i) < \ell(w)$. The reflection s_i is a descent of w if and only if α_i is an inversion of w .

Recall that for $I \subseteq [n]$ we use W_I to denote the corresponding *parabolic subgroup* of W , that is, the subgroup of $GL(V)$ generated by simple reflections s_i for $i \in I$. Note that W_I is (isomorphic to) the Weyl group of Φ_I . For $\lambda \in P$ we define the parabolic permutohedron $\Pi_I(\lambda) := \text{ConvexHull } W_I(\lambda)$ and $\Pi_I^Q(\lambda) := \Pi_I(\lambda) \cap (Q + \lambda)$. An important property of parabolic subgroups is the existence of distinguished coset representatives: each (left) coset wW_I in W contains a unique element of minimal length. We use W^I for the set of minimal length coset representatives of W_I . There is even an explicit description: $W^I := \{w \in W : s_i \text{ is not a descent of } w \text{ for all } i \in I\}$ (see for instance [BB05, §2.4]).

Recall that for any $\lambda \in P$ we use λ_{dom} to denote the dominant element of $W(\lambda)$. And for a dominant weight $\lambda = \sum_{i=1}^n c_i \omega_i \in P_{\geq 0}$, we define $I_\lambda^0 := \{i \in [n] : c_i = 0\}$.

Proposition 6.1. *For $\lambda \in P_{\geq 0}$, the stabilizer of λ in W is $W_{I_\lambda^0}$.*

Proof. This (straightforward proposition) is [Hum72, Lemma 10.2B]. \square


 FIGURE 8. A graphical depiction of the piecewise-linear map $\eta_{\mathbf{k}}$.

Corollary 6.2. *For any $\lambda \in P$, $\{w \in W : w^{-1}(\lambda) \in P_{\geq 0}\}$ is a coset of $W_{I_{\lambda_{\text{dom}}}}^0$.*

Proof. First let us show that if $w^{-1}(\lambda)$ is dominant then $(ww')^{-1}(\lambda)$ is dominant for any $w' \in W_{I_{\lambda_{\text{dom}}}}^0$. This is clear: $(ww')^{-1}(\lambda) = (w')^{-1}(w^{-1}(\lambda)) = (w')^{-1}(\lambda_{\text{dom}}) = \lambda_{\text{dom}}$ since w' is in the stabilizer of λ_{dom} by Proposition 6.1. Next let us show that if $w^{-1}(\lambda)$ is dominant and $(w')^{-1}(\lambda)$ is dominant then $w' = ww''$ for some $w'' \in W_{I_{\lambda_{\text{dom}}}}^0$. This is also clear: $w^{-1}(w'(\lambda_{\text{dom}})) = w^{-1}(\lambda) = \lambda_{\text{dom}}$, so $w^{-1}w'$ is in the stabilizer of λ_{dom} , that is, $w^{-1}w' = w''$ for some $w'' \in W_{I_{\lambda_{\text{dom}}}}^0$ thanks to Proposition 6.1, as claimed. \square

In light of Corollary 6.2, for $\lambda \in P$ we define w_λ to be the minimal length element of $\{w \in W : w^{-1}(\lambda) \in P_{\geq 0}\}$. Hence, for $\lambda \in P_{\geq 0}$ we have (by the Orbit-Stabilizer Theorem) that $W_{I_\lambda}^0 = \{w_\mu : \mu \in W(\lambda)\}$ and $w_\mu \neq w_{\mu'}$ for $\mu \neq \mu' \in W(\lambda)$. Another way to think about w_λ : λ may belong to the closure of many chambers, but there will be a unique chamber wC_0 with w of minimal length such that λ belongs to the closure of wC_0 and this is when $w = w_\lambda$. Then for $\mathbf{k} \in \mathbb{N}[\Phi]^W$, we define the map $\eta_{\mathbf{k}} : P \rightarrow P$ by setting $\eta_{\mathbf{k}}(\lambda) := \lambda + w_\lambda(\rho_{\mathbf{k}})$ for all $\lambda \in P$ (where, as above, we have $\rho_{\mathbf{k}} := \sum_{i=1}^n \mathbf{k}(\alpha_i)\omega_i$).

This map $\eta_{\mathbf{k}}$ will be of crucial importance for us in our investigation of both the symmetric and truncated interval-firing processes and the relationship between these two processes. Figure 8 gives a graphical depiction of $\eta_{\mathbf{k}}$: as we can see, this map “dilates” space by translating the chambers radially outwards; a point not inside any chamber travels in the same direction as the chamber closest to the fundamental chamber among those chambers whose closure the point lies in. The following proposition lists some very basic properties of $\eta_{\mathbf{k}}$.

Proposition 6.3.

- For any $\mathbf{k}, \mathbf{m} \in \mathbb{N}[\Phi]^W$, we have $\eta_{\mathbf{k}+\mathbf{m}} = \eta_{\mathbf{m}}(\eta_{\mathbf{k}})$.
- For any $\mathbf{k} \in \mathbb{N}[\Phi]^W$, the map $\eta_{\mathbf{k}} : P \rightarrow P$ is injective.

Proof. For the first bullet point: let $\lambda \in P$. Set $\lambda' := \eta_{\mathbf{k}}(\lambda) = \lambda + w_\lambda(\rho_{\mathbf{k}})$. Observe that $\lambda'_{\text{dom}} = w_\lambda^{-1}(\eta_{\mathbf{k}}(\lambda)) = \lambda_{\text{dom}} + \rho_{\mathbf{k}}$. Hence, $I_{\lambda'_{\text{dom}}}^0 \subseteq I_{\lambda_{\text{dom}}}^0$. This means the cosets of $W_{I_{\lambda_{\text{dom}}}^0}$ are unions of cosets of $W_{I_{\lambda'_{\text{dom}}}^0}$. But we just saw that $w_\lambda \in w_{\lambda'} W_{I_{\lambda'_{\text{dom}}}^0}$, because $w_\lambda^{-1}(\lambda')$ is dominant. So w_λ must be the minimal length element of $w_{\lambda'} W_{I_{\lambda'_{\text{dom}}}^0}$ (since it is the minimal length element of a superset of $w_{\lambda'} W_{I_{\lambda'_{\text{dom}}}^0}$). Hence $w_{\lambda'} = w_\lambda$. This means that $\eta_{\mathbf{m}}(\eta_{\mathbf{k}}(\lambda)) = \lambda + w_\lambda(\rho_{\mathbf{k}}) + w_\lambda(\rho_{\mathbf{m}}) = \lambda + w_\lambda(\rho_{\mathbf{k}+\mathbf{m}}) = \eta_{\mathbf{k}+\mathbf{m}}(\lambda)$ and thus the claim is proved.

For the second bullet point: suppose $\lambda, \mu \in P$ with $\eta_{\mathbf{k}}(\lambda) = \eta_{\mathbf{k}}(\mu)$. First of all, since $\eta_{\mathbf{k}}(\lambda)_{\text{dom}} = \lambda_{\text{dom}} + \rho_{\mathbf{k}}$ and similarly for μ , we have $\lambda_{\text{dom}} = \mu_{\text{dom}}$. Let $m \gg 0 \in \mathbb{Z}$ be some very large constant. From the first bullet point we know $\eta_{\mathbf{k}+m}(\lambda) = \eta_{\mathbf{k}+m}(\mu)$ and hence $\lambda + w_\lambda(\rho_{\mathbf{k}+m}) = \mu + w_\mu(\rho_{\mathbf{k}+m})$. But $\rho_{\mathbf{k}+m}$ is inside the fundamental chamber C_0 , and hence $w(\rho_{\mathbf{k}+m}) = w'(\rho_{\mathbf{k}+m})$ if and only if $w = w'$. Moreover, by taking m large enough we can guarantee that $w(\rho_{\mathbf{k}+m})$ and $w'(\rho_{\mathbf{k}+m})$ are very far away from one another for $w \neq w'$. Hence $\lambda + w_\lambda(\rho_{\mathbf{k}+m}) = \mu + w_\mu(\rho_{\mathbf{k}+m})$ in fact forces $w_\lambda = w_\mu$. But $w_\lambda = w_\mu$ together with $\lambda_{\text{dom}} = \mu_{\text{dom}}$ means $\lambda = \mu$ and thus the claim is proved. \square

In light of Proposition 6.3 it makes sense to set $\eta := \eta_1$ so that $\eta_{\mathbf{k}} = \eta^k$. Now we proceed to explain how $\eta_{\mathbf{k}}$ labels the sinks of $\Gamma_{\text{sym}, \mathbf{k}}$.

For a dominant weight $\lambda = \sum_{i=1}^n c_i \omega_i \in P_{\geq 0}$, define $I_\lambda^{0,1} := \{i \in [n] : c_i \in \{0, 1\}\}$.

Proposition 6.4. *Let $\lambda \in P$ with $\langle \lambda, \alpha^\vee \rangle \neq -1$ for all $\alpha \in \Phi^+$. Then $w_\lambda(\Phi_{I_\lambda^{0,1}}^+)$ is a subset of positive roots.*

Proof. It suffices to show that $w_\lambda(\alpha_i)$ is positive for all $i \in I_\lambda^{0,1}$. Suppose that $w_\lambda(\alpha_i)$ is negative for some $i \in I_\lambda^{0,1}$, i.e., that s_i is a descent of w_λ . Note $\langle \lambda_{\text{dom}}, \alpha_i^\vee \rangle \in \{0, 1\}$. If $\langle \lambda_{\text{dom}}, \alpha_i^\vee \rangle = 1$, then $\langle \lambda_{\text{dom}}, -\alpha_i^\vee \rangle = -1$ so $\langle \lambda, -w_\lambda(\alpha_i)^\vee \rangle = -1$, which contradicts that $\langle \lambda, \alpha^\vee \rangle \neq -1$ for all $\alpha \in \Phi^+$. But since w_λ is the minimal length representative of $w_\lambda W_{I_\lambda^{0,1}}$, it cannot have any descents s_j with $j \in I_\lambda^{0,1}$. Hence we cannot have that $\langle \lambda_{\text{dom}}, \alpha_i^\vee \rangle = 0$ either. Thus it must be that $w_\lambda(\alpha_i)$ is positive for all $i \in I_\lambda^{0,1}$. \square

Proposition 6.5. *For a dominant weight $\mu \in P_{\geq 0}$, we have that*

$$W_\mu^{0,1} = \{w_\lambda : \lambda \in P, \lambda_{\text{dom}} = \mu, \langle \lambda, \alpha^\vee \rangle \neq -1 \text{ for all } \alpha \in \Phi^+\}.$$

Proof. Let $\lambda \in P$ with $\lambda_{\text{dom}} = \mu$ and first suppose that $\langle \lambda, \alpha^\vee \rangle = -1$ for some $\alpha \in \Phi^+$. Then we have $\langle w_\lambda^{-1}(\lambda), w_\lambda^{-1}(\alpha)^\vee \rangle = -1$. But since $w_\lambda^{-1}(\lambda) = \lambda_{\text{dom}}$ is dominant, this means $w_\lambda^{-1}(\alpha)$ is a negative root; moreover, the only way $\langle \lambda_{\text{dom}}, w_\lambda^{-1}(\alpha)^\vee \rangle = -1$ is possible is if all the simple coroots α_i^\vee appearing in the expansion of $-w_\lambda^{-1}(\alpha)^\vee$ have $i \in I_\lambda^{0,1}$. This implies that $w_\lambda(\alpha_i)$ is negative for some $i \in I_\lambda^{0,1}$. But then s_i would be a descent of w_λ , and hence w_λ cannot be the minimal length element of $w_\lambda W_{I_\lambda^{0,1}}$.

If $\lambda \in P$ with $\lambda_{\text{dom}} = \mu$ satisfies $\langle \lambda, \alpha^\vee \rangle \neq -1$ for all $\alpha \in \Phi^+$, then we have seen in Proposition 6.4 that w_λ has no descents s_i with $i \in I_\lambda^{0,1}$ and hence indeed $w_\lambda \in W_\mu^{0,1}$. On the other hand, since $W_\mu^0 \subseteq W_\mu^{0,1}$, the cosets of $W_\mu^{0,1}$ are unions of cosets of W_μ^0 .

and hence the minimal length element of any coset of $W_{I_\mu^{0,1}}$ must be of the form w_λ for some $\lambda \in P$ with $\lambda_{\text{dom}} = \mu$. \square

Lemma 6.6. *For any $\mathbf{k} \in \mathbb{N}[\Phi]^W$, the sinks of $\Gamma_{\text{sym}, \mathbf{k}}$ are*

$$\{\eta_{\mathbf{k}}(\lambda) : \lambda \in P, \langle \lambda, \alpha^\vee \rangle \neq -1 \text{ for all } \alpha \in \Phi^+\}$$

Proof. First suppose that $\lambda \in P$ satisfies $\langle \lambda, \alpha^\vee \rangle \neq -1$ for all $\alpha \in \Phi^+$. Let $\alpha \in \Phi^+$. If $\alpha \in w_\lambda(\Phi_{I_{\lambda_{\text{dom}}}^{0,1}})$, then $\langle \eta_{\mathbf{k}}(\lambda), \alpha^\vee \rangle = \langle \lambda_{\text{dom}} + \rho_{\mathbf{k}}, w_\lambda^{-1}(\alpha)^\vee \rangle \geq \mathbf{k}(\alpha)$ since $w_\lambda^{-1}(\alpha) \in \Phi^+$ by Proposition 6.4. So now consider $\alpha \notin w_\lambda(\Phi_{I_{\lambda_{\text{dom}}}^{0,1}})$. Then $w_\lambda^{-1}(\alpha)$ may be positive or negative, but $|\langle \lambda_{\text{dom}}, w_\lambda(\alpha)^\vee \rangle| \geq 2$ (because λ_{dom} has an ω_i coefficient of at least 2 for some $i \notin I_{\lambda_{\text{dom}}}^{0,1}$ such that α_i^\vee appears in the expansion of $\pm w_\lambda(\alpha)^\vee$). Hence

$$|\langle \eta_{\mathbf{k}}(\lambda), \alpha^\vee \rangle| = |\langle \lambda_{\text{dom}} + \rho_{\mathbf{k}}, w_\lambda^{-1}(\alpha)^\vee \rangle| \geq \mathbf{k}(\alpha) + 2,$$

which means that $\langle \eta_{\mathbf{k}}(\lambda), \alpha^\vee \rangle \notin [-\mathbf{k}(\alpha) - 1, \mathbf{k}(\alpha) - 1]$. Thus $\eta_{\mathbf{k}}(\lambda)$ is a sink of $\Gamma_{\text{sym}, \mathbf{k}}$.

Now suppose μ is a sink of $\Gamma_{\text{sym}, \mathbf{k}}$. Since $\langle \mu, \alpha^\vee \rangle \notin [-\mathbf{k}(\alpha) - 1, \mathbf{k}(\alpha) - 1]$ for $\alpha \in \Phi^+$, in particular $|\langle \mu, \alpha^\vee \rangle| \geq \mathbf{k}(\alpha)$ for all $\alpha \in \Phi^+$. This means $\langle \mu_{\text{dom}}, \alpha^\vee \rangle \geq \mathbf{k}(\alpha)$ for all $\alpha \in \Phi^+$. Hence $\mu_{\text{dom}} = \mu' + \rho_{\mathbf{k}}$ for some dominant $\mu' \in P_{\geq 0}$. Suppose to the contrary that w_μ is not the minimal length element of $w_\mu W_{I_{\mu'}^{0,1}}$. Then there exists a descent s_i of w_μ with $i \in I_{\mu'}^{0,1}$. But then

$$\langle \mu, -w_\mu(\alpha_i)^\vee \rangle = \langle \mu_{\text{dom}}, -\alpha_i^\vee \rangle = -\langle \mu', \alpha_i^\vee \rangle - \langle \rho_{\mathbf{k}}, \alpha_i^\vee \rangle \geq -\mathbf{k}(\alpha_i) - 1,$$

and also $\langle \mu, -w_\mu(\alpha_i)^\vee \rangle = -\langle \mu_{\text{dom}}, \alpha_i^\vee \rangle \leq 0$. This would imply that μ is not a sink of $\Gamma_{\text{sym}, \mathbf{k}}$, since $-w_\mu(\alpha_i) \in \Phi^+$. So w_μ must be the minimal length element of $w_\mu W_{I_{\mu'}^{0,1}}$. Thanks to Proposition 6.5, this means $w_\mu = w_\lambda$ for some $\lambda \in P$ with $\lambda_{\text{dom}} = \mu'$ and $\langle \lambda, \alpha^\vee \rangle \neq -1$ for all $\alpha \in \Phi^+$. Moreover, $\mu = w_\mu(\mu_{\text{dom}}) = \lambda + w_\lambda(\rho_{\mathbf{k}}) = \eta_{\mathbf{k}}(\lambda)$, as claimed. \square

7. TRAVERSE LENGTHS OF PERMUTOHEDRA

Our goal will now be to describe the connected components of $\Gamma_{\text{sym}, \mathbf{k}}$, with the eventual aim of establishing confluence of $\xrightarrow{\text{sym}, \mathbf{k}}$. (By *connected component* of a directed graph, we mean a connected component of its underlying undirected graph.) We will show over the course of the next several sections that the connected components are contained in certain permutohedra; from this confluence will follow easily. First we need to discuss traverse lengths.

Definition 7.1. For a root $\alpha \in \Phi$, an α -string of length ℓ is a subset of P of the form $\{\mu, \mu - \alpha, \mu - 2\alpha, \dots, \mu - \ell\alpha\}$ for some weight $\mu \in P$. For a dominant weight $\lambda \in P_{\geq 0}$, an α -traverse in the discrete permutohedron $\Pi^Q(\lambda)$ is a maximal (as a set) α -string that belongs to $\Pi^Q(\lambda)$. Concretely, it is an α -string $\{\mu, \mu - \alpha, \mu - 2\alpha, \dots, \mu - \ell\alpha\} \subseteq \Pi^Q(\lambda)$ such that $\mu + \alpha, \mu - (\ell + 1)\alpha \notin \Pi^Q(\lambda)$. Finally, for a dominant weight $\lambda \in P_{\geq 0}$, the *traverse length* $\mathbf{l}_\lambda \in \mathbb{Z}[\Phi]^W$ is given by

$$\mathbf{l}_\lambda(\alpha) := \text{the minimal length } \ell \text{ of an } \alpha\text{-traverse in } \Pi^Q(\lambda).$$

Clearly, by the W -symmetry of permutohedra, the traverse length is W -invariant and hence really does belong to $\mathbb{Z}[\Phi]^W$.

Lemma 7.2. *For $\lambda \in P$ and $\alpha \in \Phi$, any α -traverse $\{\mu, \mu - \alpha, \dots, \mu - \ell\alpha\} \subseteq \Pi^Q(\lambda)$ is symmetric with respect to the reflection s_α , i.e., $s_\alpha(\mu - i\alpha) = \mu - (\ell - i)\alpha$ for all $i = 0, \dots, \ell$. Its length is $\ell = \langle \mu, \alpha^\vee \rangle$. In particular, $\langle \mu, \alpha^\vee \rangle \geq 0$.*

Proof. By the W -symmetry of discrete permutohedra, we have $s_\alpha(\Pi^Q(\lambda)) = \Pi^Q(\lambda)$, which implies the first sentence. The second sentence then follows from

$$\mu - \ell\alpha = s_\alpha(\mu) = \mu - \langle \mu, \alpha^\vee \rangle \alpha.$$

The last sentence is clear because the length ℓ must be nonnegative. \square

Lemma 7.2 implies the following reformulation of the definition of \mathbf{l}_λ .

Corollary 7.3. *For $\lambda \in P$, the traverse length \mathbf{l}_λ is given by*

$$\mathbf{l}_\lambda(\alpha) = \min(\{\langle \mu, \alpha^\vee \rangle : \mu \in \Pi^Q(\lambda), \mu + \alpha \notin \Pi^Q(\lambda)\}).$$

Corollary 7.3 explains the connection of traverse length to interval-firing: we are going to prove that interval-firing processes get “trapped” inside of permutohedra because the traverse lengths of these permutohedra are large (and hence if μ is inside such a permutohedron but $\mu + \alpha$ is not, $\langle \mu, \alpha^\vee \rangle$ must be so large that it is outside the fierability interval of our process). To do this we need a formula for traverse length. In most cases, the traverse length of a permutohedron in a given direction α is realized on some edge of the permutohedron in direction α . However, there are some strange exceptions to this general rule, for which we need the concept of “funny” weights.

Definition 7.4. If Φ is simply laced, then there are no funny weights. So suppose Φ is not simply laced. Then there is a unique long simple root α_l and short simple root α_s with $\langle \alpha_l, \alpha_s^\vee \rangle \neq 0$. We say the dominant weight $\lambda = \sum_{i=1}^n c_i \omega_i \in P_{\geq 0}$ is *funny* if $c_s = 0$ and $c_l \geq 1$ and $c_i \geq c_l$ for all i such that α_i is long.

Example 7.5. With the numbering of simple roots as in Figure 2, if $\Phi = B_n$ then $\lambda = \sum_{i=1}^n c_i \omega_i \in P_{\geq 0}$ is funny if $c_1, \dots, c_{n-2} \geq c_{n-1} \geq 1$ and $c_n = 0$. If $\Phi = C_n$, then λ is funny if $c_{n-1} = 0$ and $c_n \geq 1$.

For a dominant weight $\lambda = \sum_{i=1}^n c_i \omega_i \in P_{\geq 0}$, define $\mathbf{m}_\lambda \in \mathbb{Z}[\Phi]^W$ by setting

$$\mathbf{m}_\lambda(\alpha) := \min(\{c_i : \alpha \in W(\alpha_i)\}).$$

Theorem 7.6. *For a dominant weight $\lambda \in P_{\geq 0}$, we have*

$$\mathbf{l}_\lambda(\alpha) = \begin{cases} \mathbf{m}_\lambda(\alpha) - 1 & \text{if } \alpha \text{ is long and } \lambda \text{ is funny,} \\ \mathbf{m}_\lambda(\alpha) & \text{otherwise.} \end{cases}$$

Proof. Let $\lambda = \sum_{i=1}^n c_i \omega_i \in P_{\geq 0}$. The α_i -traverse $\{\lambda, \lambda - \alpha, \dots, \lambda - \ell\alpha = s_i(\lambda)\}$, which is contained in the edge $[\lambda, s_i(\lambda)]$ of the permutohedron $\Pi(\lambda)$, has length equal to $\ell = \langle \lambda, \alpha_i^\vee \rangle = c_i$. By the W -symmetry of the traverse length (and because any root is W -conjugate to some simple root), it follows that $\mathbf{l}_\lambda \leq \mathbf{m}_\lambda$.

We will show that in most of the cases (except the case with long roots and funny weights) we actually have $\mathbf{l}_\lambda = \mathbf{m}_\lambda$. We need to show that the length of any α -traverse

in $\Pi^Q(\lambda)$ is greater than or equal to $\mathbf{m}_\lambda(\alpha)$, i.e., for $\mu \in \Pi^Q(\lambda)$ such that $\mu + \alpha \notin \Pi^Q(\lambda)$, we have $\langle \mu, \alpha^\vee \rangle \geq \mathbf{m}_\lambda(\alpha)$.

If $\mathbf{m}_\lambda(\alpha) = 0$, then we automatically get $\mathbf{l}_\lambda(\alpha) = \mathbf{m}_\lambda(\alpha) = 0$, because $\mathbf{l}_\lambda(\alpha)$. So let us assume that $\mathbf{m}_\lambda(\alpha) \geq 1$.

Let $\mu \in \Pi^Q(\lambda)$ be such that $\mu + \alpha \notin \Pi^Q(\lambda)$. Since $\mu + \alpha \in Q + \lambda$, we deduce that $\mu + \alpha \notin \Pi(\lambda)$. This means that the line segment $[\mu, \mu + \alpha]$ must “exit” the permutohedron $\Pi(\lambda)$ at some point $v \in V$, i.e., there exists a unique point $v = \mu + t\alpha$, where $t \in \mathbb{R}$, with $v \in \Pi(\lambda)$ but $\mu + q\alpha \notin \Pi(\lambda)$ for any $q > t$. We have $0 \leq t < 1$.

Let F be the minimal (by inclusion) face of $\Pi(\lambda)$ that contains the point v . For any vertex ν of the face F , we have $\nu \in \Pi(\lambda)$ and $\nu + \alpha \notin \Pi(\lambda)$.

The minimal value of the linear form $\langle \cdot, \alpha^\vee \rangle$ on the face F should be reached at a vertex ν of F . By the W -symmetry of $\Pi(\lambda)$, we assume without loss of generality that this minimum is achieved at $\nu = \lambda$. So we have $\langle \lambda, \alpha^\vee \rangle \leq \langle v, \alpha^\vee \rangle$.

Let $\alpha = \sum_{i=1}^n a_i \alpha_i$, where the a_i are either all nonnegative or all nonpositive. Then we have $\alpha^\vee = \sum_{i=1}^n \tilde{a}_i \alpha_i$ where $\tilde{a}_i = \frac{\langle \alpha_i, \alpha_i \rangle}{\langle \alpha, \alpha \rangle} a_i$.

Any root α is W -conjugate to at least one simple root that appears with nonzero coefficient in its expansion in terms of the simple roots. So there exists an index j such that $\alpha_j \in W(\alpha)$ and $a_j = \tilde{a}_j \neq 0$. We have $c_j \geq \mathbf{m}_\lambda(\alpha) \geq 1$.

We have $\lambda \in \Pi(\lambda)$ and $\lambda + \alpha \notin \Pi(\lambda)$. So $\langle \lambda, \alpha^\vee \rangle \geq 0$, because $\langle \lambda, \alpha^\vee \rangle$ is the length of the α -traverse that starts at λ , which is always nonnegative. Therefore we have $\langle \lambda, \alpha^\vee \rangle = \sum_{i=1}^n \tilde{a}_i c_i \geq 0$; moreover, all nonzero terms in this expression have the same sign and at least one term $\tilde{a}_j c_j$ is nonzero. It follows that $a_1, \dots, a_n \geq 0$, i.e., α is a positive root.

If λ is strictly in the fundamental chamber, then all edges of $\Pi(\lambda)$ coming out of λ must be in the direction of a negative simple root. This is not true for general $\lambda \in P_{\geq 0}$, but the edges of $\Pi(\lambda)$ coming out of λ that are not in the direction of a negative simple root must immediately leave the dominant chamber. Hence if we let $x \in V$ be some generic point in the interior of the face F very close to λ , by acting by $W_{I_\lambda^0}$ we can transport x to the dominant chamber while fixing λ . Thus, we may assume that the affine span of F is spanned by simple roots. So let $I \subseteq [n]$ be the minimal set of indices such that the face F belongs to the affine subspace $\lambda + \text{Span}_{\mathbb{R}}(\{\alpha_i : i \in I\})$.

We have $\mu = v - t\alpha = (\lambda - \sum_{i \in I} b_i \alpha_i) - t\alpha$ for real numbers $0 \leq t < 1$ and $b_i \geq 0$, $i \in I$. Thus $\langle \mu, \alpha^\vee \rangle = \langle v, \alpha^\vee \rangle - t\langle \alpha, \alpha^\vee \rangle = \langle v, \alpha^\vee \rangle - 2t \geq \langle \lambda, \alpha^\vee \rangle - 2t > \langle \lambda, \alpha^\vee \rangle - 2$. Moreover, since both $\langle \mu, \alpha^\vee \rangle$ and $\langle \lambda, \alpha^\vee \rangle - 2$ are integers, and the first is strictly greater than the second, we get

$$\langle \mu, \alpha^\vee \rangle \geq \langle \lambda, \alpha^\vee \rangle - 1 = \left(\sum_{i=1}^n \tilde{a}_i c_i \right) - 1.$$

We already noted that the last expression involves at least one nonzero term $\tilde{a}_j c_j$ such that $\alpha_j \in W(\alpha)$. So $\tilde{a}_j c_j \geq c_j \geq \mathbf{m}_\lambda(\alpha)$ and thus $\langle \mu, \alpha^\vee \rangle \geq \mathbf{m}_\lambda(\alpha) - 1$.

We need to prove just a slightly stronger inequality $\langle \mu, \alpha^\vee \rangle \geq \mathbf{m}_\lambda(\alpha)$.

If $\sum_{\alpha_i \in W(\alpha)} \tilde{a}_i \geq 2$, we get

$$\langle \mu, \alpha^\vee \rangle \geq \sum_{i=1}^n \tilde{a}_i c_i - 1 \geq \sum_{\alpha_i \in W(\alpha)} \tilde{a}_i c_i \geq 2\mathbf{m}_\lambda(\alpha) - 1 \geq \mathbf{m}_\lambda(\alpha),$$

as needed. So we now assume that $\sum_{\alpha_i \in W(\alpha)} \tilde{a}_i = 1$.

Since α does not belong to the subspace spanned by the α_i for $i \in I$, there is $r \notin I$ such that $a_r \geq 1$.

If $a_r = 1$, then, from the fact that $\lambda - \mu = (\sum_{i \in I} b_i \alpha_i) + t\alpha$ belongs to the root lattice Q and thus is an *integer* linear combination of the simple roots, we deduce that in fact $t \in \mathbb{Z}$ and thus $t = 0$. In this case get $\langle \mu, \alpha^\vee \rangle \geq \langle \lambda, \alpha^\vee \rangle \geq a_j c_j \geq \mathbf{m}_\lambda(\alpha)$, as needed. So we now assume that $c_r \geq 2$.

Then note that $\alpha_r \notin W(\alpha)$, because we assumed $\sum_{\alpha_i \in W(\alpha)} \tilde{a}_i = \sum_{\alpha_i \in W(\alpha)} a_i = 1$.

If there is an index q such that $a_q \notin W(\alpha)$, $\tilde{a}_q \geq 1$ and $c_q \geq 1$, we have

$$\langle \mu, \alpha^\vee \rangle = \left(\sum_{i=1}^n \tilde{a}_i c_i \right) - 1 \geq \tilde{a}_j c_j + \tilde{a}_q c_q - 1 \geq \tilde{a}_j c_j \geq \mathbf{m}_\lambda(\alpha),$$

as needed.

The only possibility which is not covered by the above discussion is when:

- (1) There is exactly one nonzero term $a_j \alpha_j$ in the expansion $\alpha = \sum_{i=1}^n a_i \alpha_i$ such that $\alpha_j \in W(\alpha)$. For this term, $a_j = 1$ and $c_j = \mathbf{m}_\lambda(\alpha) \geq 1$.
- (2) There is at least one more nonzero term $a_i \alpha_i$ in that expansion. For all such terms, $\alpha_i \notin W(\alpha)$, $a_i \geq 2$, and $c_i = 0$.

We claim that these conditions imply that α is a long root. This is easy to check by hand for $\Phi = B_n, C_n$, or G_2 . One does not need to check Type F_4 separately, because in this case there are two long simple roots and two short simple roots, but the expansion of α involves either only one short simple root or only one long simple root. We leave it as an exercise for the reader to find a uniform root theoretic argument of the fact that conditions (1) and (2) above imply that α is long.

Also, we claim that conditions (1) and (2) above imply that λ is a funny weight. Indeed, it is a well-known and simple fact that for any root $\alpha = \sum_{i=1}^n a_i \alpha_i$, the set of $i \in [n]$ for which $a_i \neq 0$ must be a connected subset of the Dynkin diagram. Hence indeed the α_j in condition (1) must be the long simple root α_l , and one of the α_i in condition (2) must be the short simple root α_s (with notation as in Definition 7.4). Note also that $\mathbf{m}_\lambda(\alpha) = c_l$ forces $c_i \geq c_l$ for all i such that α_i is long.

In this “long and funny” case we can only get the (slightly) weaker inequality:

$$\langle \mu, \alpha^\vee \rangle \geq \mathbf{m}_\lambda(\alpha) - 1.$$

It remains to show that this last inequality is tight in this “long and funny” case. Let us concentrate on the 2-dimensional face of the permutohedron $\Pi(\lambda)$ contained in the affine subspace $\lambda + \text{Span}_{\mathbb{R}}(\{\alpha_l, \alpha_s\})$ (with notation as in Definition 7.4).

This face is equivalent to the 2-dimensional W' -permutohedron $\Pi_{W'}(\lambda')$ corresponding to the sub-root system Φ' of rank 2 with simple roots α_l and α_s , and fundamental




FIGURE 9. The graph $\Gamma_{\text{sym},\mathbf{k}}$ in Example 8.3. The permutohedron $\Pi(\rho_{\mathbf{k}})$ is shown in red.

weights ω'_1 (corresponding to α_l) and ω'_2 (corresponding to α_s), where W' is the Weyl group of Φ' , and $\lambda' = \mathbf{m}_\lambda(\alpha)\omega'_1 + 0\omega'_2$.

The 2-dimensional root system Φ' which must be equal to either B_2 or G_2 . In this situation there in fact is a $\mu \in \Pi_{W'}^Q(\lambda')$ with $\mu + \alpha \notin \Pi_{W'}^Q(\lambda')$ for some long $\alpha \in \Phi'$ such that $\langle \mu, \alpha^\vee \rangle = \mathbf{m}_\lambda(\alpha) - 1$: indeed, we can take $\alpha := \alpha_l$ and $\mu := (\mathbf{m}_\lambda(\alpha) - 1)\omega'_1$ for B_2 or $\alpha := \alpha_l$ and $\mu := (\mathbf{m}_\lambda(\alpha) - 1)\omega'_1 + \omega'_2$ for G_2 .

This finishes the proof of the theorem. \square

8. THE PERMUTOHEDRON NON-ESCAPING LEMMA

We need to place some restrictions on our parameter \mathbf{k} so that funny weights do not occur in our analysis of the relevant permutohedra traverse lengths. For this we have the notion of “goodness.”

Definition 8.1. If Φ is simply laced, then every $\mathbf{k} \in \mathbb{N}[\Phi]^W$ is good. So suppose Φ is not simply laced and let $\mathbf{k} \in \mathbb{N}[\Phi]^W$. Then there exist $k_s, k_l \in \mathbb{Z}$ with $\mathbf{k}(\alpha) = k_s$ if α is short and $\mathbf{k}(\alpha) = k_l$ if α is long. We say \mathbf{k} is *good* if $k_s = 0 \Rightarrow k_l = 0$. Note in particular that if $\mathbf{k} = k \geq 0$ is constant, then it is good.

Now we can prove the following permutohedron non-escaping lemma, which says that certain discrete permutohedra “trap” the symmetric interval-firing process inside of them.

Lemma 8.2. *Let $\mathbf{k} \in \mathbb{N}[\Phi]^W$ be good and let $\Gamma := \Gamma_{\text{sym},\mathbf{k}}^{\text{un}}$. Let $\lambda \in P_{\geq 0}$. Then there is no directed edge (μ, μ') in Γ with $\mu \in \Pi^Q(\eta_{\mathbf{k}}(\lambda))$ and $\mu' \notin \Pi^Q(\eta_{\mathbf{k}}(\lambda))$.*

Proof. First suppose Φ is not simply laced and $k_s = 0$. Then also $k_l = 0$, i.e., $\mathbf{k} = 0$, since \mathbf{k} is good. Hence $\rho_{\mathbf{k}} = 0$, so $\eta_{\mathbf{k}}(\lambda) = \lambda$. If $\mu \in \Pi^Q(\lambda)$ but $\mu + \alpha \notin \Pi^Q(\lambda)$, then by Corollary 7.3 we have $\langle \mu, \alpha^\vee \rangle \geq \mathbf{I}_\lambda(\alpha)$. Note that by definition $\mathbf{I}_\lambda(\alpha) \geq 0$. But this means $\langle \mu, \alpha^\vee \rangle \geq \mathbf{k}(\alpha)$, so indeed $(\mu, \mu + \alpha)$ cannot be a directed edge of Γ .

Now suppose either Φ is simply laced or Φ is not simply laced but $k_s \geq 1$. Then note that $\rho_{\mathbf{k}}$ is not funny. Hence by Theorem 7.6 we conclude that $\mathbf{I}_{\eta_{\mathbf{k}}(\lambda)}(\alpha) \geq \mathbf{k}(\alpha)$. If $\mu \in \Pi^Q(\eta_{\mathbf{k}}(\lambda))$ but $\mu + \alpha \notin \Pi^Q(\eta_{\mathbf{k}}(\lambda))$, then $\langle \mu, \alpha^\vee \rangle \geq \mathbf{I}_{\eta_{\mathbf{k}}(\lambda)}(\alpha)$ by Corollary 7.3. This means $\langle \mu, \alpha^\vee \rangle \geq \mathbf{k}(\alpha)$, so indeed $(\mu, \mu + \alpha)$ cannot be a directed edge of Γ . \square

Example 8.3. Lemma 8.2 is false in general without the goodness assumption. Indeed, suppose $\Phi = B_2$ and $\mathbf{k} \in \mathbb{N}[\Phi]^W$ is given by $k_s := 0$ and $k_l := 1$. Then Figure 9 depicts (a portion of) the graph $\Gamma_{\text{sym}, \mathbf{k}}$. In this picture we only show elements of the root lattice Q . The permutohedron $\Pi(\rho_{\mathbf{k}}) = \Pi(\omega_1)$ is shown in red. Observe that although $0 \in \Pi^Q(\rho_{\mathbf{k}})$ and $\alpha_1 \notin \Pi^Q(\rho_{\mathbf{k}})$, we have an edge $(0, \alpha_1)$ in $\Gamma_{\text{sym}, \mathbf{k}}$.

We also need a “lower-dimensional” version of the permutohedron non-escaping lemma that says that these interval-firing processes get trapped inside of permutohedra of parabolic subgroups of W . This is established in the following lemma and corollary.

Lemma 8.4. *Let $\mathbf{k} \in \mathbb{N}[\Phi]^W$ and $\Gamma := \Gamma_{\text{sym}, \mathbf{k}}^{\text{un}}$. Let $\lambda \in P_{\geq 0}$. Then if $(\mu, \mu + \alpha)$ is a directed edge in Γ with $\mu \in \Pi_{I_\lambda}^Q(\eta_{\mathbf{k}}(\lambda))$, we have $\alpha \in \Phi_{I_\lambda}^{0,1}$.*

Proof. Write $\eta_{\mathbf{k}}(\lambda) = \sum_{i=1}^n c_i \omega_i$. Assume to the contrary that there exists an edge $(\mu, \mu + \alpha)$ in Γ such that $\mu \in \Pi_{I_\lambda}^Q(\eta_{\mathbf{k}}(\lambda))$ but α does not belong to $\text{Span}_{\mathbb{R}}(\{\alpha_i : i \in I\})$.

Note that α is a root (positive or negative) with $-\mathbf{k}(\alpha) - 1 \leq \langle \mu, \alpha^\vee \rangle \leq \mathbf{k}(\alpha) + 1$. Let $\beta = \pm \alpha \in \Phi^+$ be the positive root. Then $\langle \mu, \beta^\vee \rangle \leq \mathbf{k}(\alpha) + 1 \leq \mathbf{k}(\beta) + 1$.

Since the point μ belongs to $\Pi_{I_\lambda}^Q(\eta_{\mathbf{k}}(\lambda))$, we deduce that the same inequality $\langle \nu, \beta^\vee \rangle \leq \mathbf{k}(\beta) + 1$ holds for some vertex ν of $\Pi_{I_\lambda}^Q(\eta_{\mathbf{k}}(\lambda))$. We have $\nu = w(\eta_{\mathbf{k}}(\lambda))$ where $w \in W_{I_\lambda}^{0,1}$. Hence we have that $\langle w(\lambda), \beta^\vee \rangle = \langle \lambda, w^{-1}(\beta)^\vee \rangle \leq \mathbf{k}(\beta) + 1$ for some $w \in W_{I_\lambda}^{0,1}$.

The action of the parabolic subgroup $W_{I_\lambda}^{0,1}$ on β^\vee does not change the coefficients b_j of the expansion $\beta^\vee = \sum_{i=1}^n b_i \alpha_i^\vee$ for all $j \notin I$, and at least one of these coefficients b_j should be strictly positive (because β^\vee is a positive coroot that does not belong to $\text{Span}_{\mathbb{R}}(\{\alpha_i : i \in I\})$). So the expansion $w^{-1}(\beta)^\vee = \sum_{i=1}^n b'_i \alpha_i^\vee$ contains some strictly positive coefficient, which means that $w^{-1}(\beta)^\vee$ is a positive coroot and thus we have $b'_i \geq 0$ for all i .

Moreover, any coroot is W -conjugate to some simple coroot that appears in its expansion with nonzero coefficient. These observations mean that we can find $j \notin I$ such that $b'_j = b_j \geq 1$, and also (possibly the same) i such that $b'_i \geq 1$ and $\alpha_i \in W(\alpha)$. Note that for this i we have $\mathbf{k}(\alpha_i) = \mathbf{k}(w^{-1}(\beta)) = \mathbf{k}(\beta) = \mathbf{k}(\alpha)$.

If $i = j$, we get $\langle \lambda, w^{-1}(\beta)^\vee \rangle \geq \langle \lambda, b'_j \alpha_j^\vee \rangle \geq \langle \lambda, \alpha_j^\vee \rangle = c_j \geq \mathbf{k}(\alpha_j) + 2 = \mathbf{k}(\alpha) + 2$ (because for $j \notin I$, $c_j \geq \mathbf{k}(\alpha_j) + 2$). But this contradicts $\langle \lambda, w^{-1}(\beta)^\vee \rangle \leq \mathbf{k}(\alpha) + 1$.

On the other hand, if $i \neq j$, we get

$$\begin{aligned} \langle \lambda, w^{-1}(\beta)^\vee \rangle &\geq \langle \lambda, b'_i \alpha_i^\vee + b'_j \alpha_j^\vee \rangle \geq \langle \lambda, \alpha_i^\vee \rangle + \langle \lambda, \alpha_j^\vee \rangle = c_i + c_j \\ &\geq \mathbf{k}(\alpha_i) + (\mathbf{k}(\alpha_j) + 2) \geq \mathbf{k}(\alpha_i) + 2 = \mathbf{k}(\alpha) + 2. \end{aligned}$$

Again, we get a contradiction. \square

Corollary 8.5. *Let $\mathbf{k} \in \mathbb{N}[\Phi]^W$ be good and $\Gamma := \Gamma_{\text{sym}, \mathbf{k}}^{\text{un}}$. Let $\lambda \in P_{\geq 0}$. Then there is no directed edge (μ, μ') in Γ with $\mu \in \Pi_{I_\lambda}^Q(\eta_{\mathbf{k}}(\lambda))$ and $\mu' \notin \Pi_{I_\lambda}^Q(\eta_{\mathbf{k}}(\lambda))$.*

Proof. This follows by combining Lemmas 8.2 and 8.4: if we have a directed edge $(\mu, \mu + \alpha)$ with $\mu \in \Pi_{I_\lambda^{0,1}}^Q(\eta_{\mathbf{k}}(\lambda))$, then $\alpha \in \Phi_{I_\lambda^{0,1}}$ by Lemma 8.4; hence this firing move is equivalent (via projection) to the same move for the sub-root system $\Phi_{I_\lambda^{0,1}}$; so by Lemma 8.2 applied to that sub-root system, we have $\mu + \alpha \in \Pi_{I_\lambda^{0,1}}^Q(\eta_{\mathbf{k}}(\lambda))$. \square

9. CONFLUENCE OF SYMMETRIC INTERVAL-FIRING

Now, as promised, we are ready to show that connected components of $\Gamma_{\text{sym}, \mathbf{k}}$ are contained inside permutohedra.

Theorem 9.1. *Let $\mathbf{k} \in \mathbb{N}[\Phi]^W$ be good. Let $\lambda \in P$ with $\langle \lambda, \alpha^\vee \rangle \neq -1$ for all $\alpha \in \Phi^+$. Let $Y_\lambda := \{\mu \in P : \mu \xleftarrow[\text{sym}, \mathbf{k}]{*} \eta_{\mathbf{k}}(\lambda)\}$ be the connected component of $\Gamma_{\text{sym}, \mathbf{k}}$ containing the sink $\eta_{\mathbf{k}}(\lambda)$. Then Y_λ is contained in $w_\lambda \Pi_{I_{\lambda_{\text{dom}}}^{0,1}}^Q(\eta_{\mathbf{k}}(\lambda_{\text{dom}}))$.*

Proof. First suppose that λ is dominant. By Corollary 8.5 there is no edge (μ, μ') in $\Gamma_{\text{sym}, \mathbf{k}}$ where one of μ, μ' is in $\Pi_{I_\lambda^{0,1}}^Q(\eta_{\mathbf{k}}(\lambda))$ and the other is not, which implies that Y_λ is contained in $\Pi_{I_\lambda^{0,1}}^Q(\eta_{\mathbf{k}}(\lambda))$. Now suppose λ is not dominant. By the preceding argument, the result is true for λ_{dom} . But then we have $Y_\lambda = w_\lambda Y_{\lambda_{\text{dom}}}$ by the W -symmetry of $\Gamma_{\text{sym}, \mathbf{k}}^{\text{un}}$, i.e., by Theorem 5.1. \square

And now we can prove half of Theorem 1.1.

Corollary 9.2. *Let $\mathbf{k} \in \mathbb{N}[\Phi]^W$ be good. Then $\xrightarrow[\text{sym}, \mathbf{k}]{} is confluent (and terminating).$*

Proof. We already saw in Proposition 4.3 that $\xrightarrow[\text{sym}, \mathbf{k}]{} is terminating. Thus, every connected component of $\Gamma_{\text{sym}, \mathbf{k}}$ contains at least one sink, and $\xrightarrow[\text{sym}, \mathbf{k}]{} is confluent as long as every connected component contains a unique sink.$$

So suppose that two sinks belong to the same connected component of $\Gamma_{\text{sym}, \mathbf{k}}$. By Lemma 6.6, we know that these sinks must be of the form $\eta_{\mathbf{k}}(\lambda)$ and $\eta_{\mathbf{k}}(\lambda')$ for $\lambda, \lambda' \in P$ with $\langle \lambda, \alpha^\vee \rangle \neq -1$ and $\langle \lambda', \alpha^\vee \rangle \neq -1$ for all $\alpha \in \Phi^+$.

By Theorem 9.1, $\eta_{\mathbf{k}}(\lambda) \in w_{\lambda'} \Pi_{I_{\lambda'_{\text{dom}}}^{0,1}}^Q(\eta_{\mathbf{k}}(\lambda'_{\text{dom}}))$ and vice-versa. In particular we have that $\eta_{\mathbf{k}}(\lambda_{\text{dom}}) \in \Pi^Q(\eta_{\mathbf{k}}(\lambda'_{\text{dom}}))$ and $\eta_{\mathbf{k}}(\lambda'_{\text{dom}}) \in \Pi^Q(\eta_{\mathbf{k}}(\lambda_{\text{dom}}))$. Proposition 2.2 then says that $\eta_{\mathbf{k}}(\lambda_{\text{dom}}) - \eta_{\mathbf{k}}(\lambda'_{\text{dom}})$ and $\eta_{\mathbf{k}}(\lambda'_{\text{dom}}) - \eta_{\mathbf{k}}(\lambda_{\text{dom}})$ are both in $Q_{\geq 0}$, which is possible only if $\eta_{\mathbf{k}}(\lambda_{\text{dom}}) = \eta_{\mathbf{k}}(\lambda'_{\text{dom}})$. That is, thanks to the injectivity of $\eta_{\mathbf{k}}$ established in Proposition 6.3, we must have $\lambda_{\text{dom}} = \lambda'_{\text{dom}}$.

But then the fact that $\eta_{\mathbf{k}}(\lambda) \in w_{\lambda'} \Pi_{I_{\lambda_{\text{dom}}}^{0,1}}^Q(\eta_{\mathbf{k}}(\lambda_{\text{dom}}))$ means that $\eta_{\mathbf{k}}(\lambda)$ is a vertex of $w_{\lambda'} \Pi_{I_{\lambda_{\text{dom}}}^{0,1}}^Q(\eta_{\mathbf{k}}(\lambda_{\text{dom}}))$, i.e., $\eta_{\mathbf{k}}(\lambda) = w_{\lambda'} w(\eta_{\mathbf{k}}(\lambda_{\text{dom}}))$ for some $w \in W_{I_{\lambda_{\text{dom}}}^{0,1}}$. Note that $(w_{\lambda'} w)^{-1}(\eta_{\mathbf{k}}(\lambda))$ is dominant. We have seen in the the proof of Proposition 6.3 that this means $(w_{\lambda'} w)^{-1}(\lambda)$ is dominant as well, or in other words, that $w_{\lambda'} w = w_\lambda w'$ for some $w' \in W_{I_{\lambda_{\text{dom}}}^0}$. This shows that $w_\lambda \in w_{\lambda'} W_{I_{\lambda_{\text{dom}}}^{0,1}}$. By Proposition 6.5, w_λ and $w_{\lambda'}$

must both be the minimal length elements of the cosets of $W_{I_\lambda^{0,1}}$ they belong to. So $w_\lambda = w_{\lambda'}$. That $\lambda_{\text{dom}} = \lambda'_{\text{dom}}$ and $w_\lambda = w_{\lambda'}$ implies that $\lambda = \lambda'$, and consequently that $\eta_{\mathbf{k}}(\lambda) = \eta_{\mathbf{k}}(\lambda')$, as required. \square

Remark 9.3. As far as we know, Theorem 9.1 and Corollary 9.2 may be true even in the case where \mathbf{k} is not good. Indeed, it appears that $\xrightarrow{\text{sym}, \mathbf{k}}$ is confluent for

all $\mathbf{k} \in \mathbb{N}[\Phi]^W$ and to prove this it would be sufficient, thanks to the diamond lemma (Lemma 3.1), to prove it for root systems of rank 2, of which there are only four: $A_1 \times A_1$, A_2 , B_2 , G_2 . All \mathbf{k} are good for simply laced root systems, so in fact one would need only check B_2 and G_2 .

10. FULL-DIMENSIONAL COMPONENTS, SATURATED COMPONENTS, AND CARTAN MATRIX CHIP-FIRING AS A LIMIT

Let $\mathbf{k} \in \mathbb{N}[\Phi]^W$ be good. For $\lambda \in P$, recall the notation $Y_\lambda := \{\mu \in P : \mu \xleftarrow{\text{sym}, \mathbf{k}}^* \eta_{\mathbf{k}}(\lambda)\}$ for the connected component of $\Gamma_{\text{sym}, \mathbf{k}}$ containing the sink $\eta_{\mathbf{k}}(\lambda)$ from the last section. By the results of the last section, all these components are distinct. In this section, we take a moment to highlight certain special components Y_λ , namely:

- those which are *full-dimensional* in the sense that their affine hulls are the whole vector space: $\text{AffineHull } Y_\lambda = V$;
- those which are full-dimensional and *saturated* in the sense that they contain all lattice points in their convex hulls: $Y_\lambda = (\text{ConvexHull } Y_\lambda) \cap (Q + \eta_{\mathbf{k}}(\lambda))$.

For the full-dimensional components: by a result we will prove later (Corollary 14.2), we have that Y_λ always contains $W(\eta_{\mathbf{k}}(\lambda_{\text{dom}}))$ for $\lambda \in P_{\geq 0}$ with $I_\lambda^{0,1} = [n]$. Hence by Theorem 9.1 we see that the full-dimensional connected components of $\Gamma_{\text{sym}, \mathbf{k}}$ are exactly Y_λ for $\lambda \in P_{\geq 0}$ with $I_\lambda^{0,1} = [n]$, i.e., those $\lambda = \sum_{i=1}^n c_i \omega_i \in P$ with $c_i \in \{0, 1\}$ for all $i \in [n]$. Clearly there are 2^n such full-dimensional components.

For the full-dimensional and saturated components: by that same Corollary 14.2, we see that Y_λ being full-dimensional and saturated is equivalent to having this component satisfy $Y_\lambda = \Pi^Q(\eta_{\mathbf{k}}(\lambda_{\text{dom}}))$. And then we have the following:

Proposition 10.1. *Let $\mathbf{k} \in \mathbb{N}[\Phi]^W$ be good. Let $\lambda \in P$ be a weight with $\langle \lambda, \alpha^\vee \rangle \neq -1$ for all $\alpha \in \Phi^+$. Let $Y_\lambda := \{\mu \in P : \mu \xleftarrow{\text{sym}, \mathbf{k}}^* \eta_{\mathbf{k}}(\lambda)\}$ be the connected component of $\Gamma_{\text{sym}, \mathbf{k}}$ containing the sink $\eta_{\mathbf{k}}(\lambda)$. Then Y_λ is equal to $\Pi^Q(\eta_{\mathbf{k}}(\lambda))$ if and only if $\lambda \in \Omega_m^0$.*

Proof. First note that if λ is a sink of $\Gamma_{\text{sym}, \mathbf{k}}$ then so is λ_{dom} and by the confluence of $\xrightarrow{\text{sym}, \mathbf{k}}$ there cannot be two sinks in a single connected component of $\Gamma_{\text{sym}, \mathbf{k}}$, so it

suffices to prove this proposition for dominant $\lambda \in P_{\geq 0}$ with $I_\lambda^{0,1} = [n]$. (Observe that if $\lambda \in \Omega_m^0$ then certainly it is of this form.)

By the polytopal characterization of minuscule weights, there exists a dominant weight $\mu \in P_{\geq 0}$ with $\mu \in \Pi^Q(\lambda)$ but $\mu \neq \lambda$ if and only if $\lambda \notin \Omega_m^0$. Hence by Proposition 2.2 there exists $\mu \in P_{\geq 0}$ with $\eta_{\mathbf{k}}(\mu) \in \Pi^Q(\eta_{\mathbf{k}}(\lambda))$ but $\eta_{\mathbf{k}}(\mu) \neq \eta_{\mathbf{k}}(\lambda)$ if and only if $\lambda \notin \Omega_m^0$. By applying W , we see that there is a sink $\eta_{\mathbf{k}}(\mu)$ of $\Gamma_{\text{sym}, \mathbf{k}}$

with $\eta_{\mathbf{k}}(\mu) \in \Pi^Q(\eta_{\mathbf{k}}(\lambda))$ but $\eta_{\mathbf{k}}(\mu) \notin W(\eta_{\mathbf{k}}(\lambda))$ if and only if $\lambda \notin \Omega_m^0$. Finally, by the permutohedron non-escaping lemma, Lemma 8.2, this means precisely that $\Pi^Q(\eta_{\mathbf{k}}(\lambda))$ is its own connected component if and only if $\lambda \in \Omega_m^0$. \square

Remark 10.2. Proposition 10.1 fails when \mathbf{k} is not good, as can be seen in Example 8.3 above: in this example, $0 \in \Pi^Q(\rho_{\mathbf{k}})$ but 0 does not belong to the connected component of $\Gamma_{\text{sym}, \mathbf{k}}$ containing $\rho_{\mathbf{k}} = \eta_{\mathbf{k}}(0) = \omega_1$.

So we see that the full-dimensional and saturated components of $\Gamma_{\text{sym}, \mathbf{k}}$ are exactly the Y_ω for $\omega \in \Omega_m^0$. There are f of these, where we recall that $f := \#P/Q$ is the index of connection of Φ . In some sense P/Q is the “sandpile group” in our setting, and in fact we have that $P/Q \simeq \text{coker}(\mathbf{C}^t)$, where \mathbf{C} is the Cartan matrix of Φ . Hence, these full-dimensional and saturated components suggest that interval-firing may possibly be connected to Cartan matrix chip-firing. The next remark explains that indeed there is some connection.

Remark 10.3. Let us explain how Cartan matrix chip-firing (which, as mentioned, has been investigated by Benkart-Klivans-Reiner [BKR18]) can be realized as a certain “limit” of symmetric interval-firing. Note that a Cartan matrix is always an M-matrix (see [BKR18, Proposition 4.1]). By associating to each vector $c = (c_1, \dots, c_n) \in \mathbb{Z}^n$ the weight $\lambda = \sum_{i=1}^n c_i \omega_i \in P$, we can view Cartan matrix chip-firing as the relation $\xrightarrow{\mathbf{C}}$ on P defined by

$$\lambda \xrightarrow{\mathbf{C}} \lambda - \alpha_i \text{ for } \lambda \in P \text{ and simple root } \alpha_i, i \in [n] \text{ with } \langle \lambda, \alpha_i^\vee \rangle \geq 2.$$

For $\lambda = \sum_{i=1}^n c_i \omega_i \in P$ and $k \in \mathbb{Z}_{\geq 0}$ set $B_k(\lambda) := \{\sum_{i=1}^n c'_i \omega_i \in P : \sum_{i=1}^n |c_i - c'_i| \leq k\}$. In other words, $B_k(\lambda)$ consists of those μ which are within distance k of λ in the “taxicab distance” on P . Note that for all $\lambda \in B_k(\rho_k)$, we have that $\langle \lambda, \alpha^\vee \rangle \geq k$ if $\alpha \in \Phi^+$ is not a simple root. In other words, for $\lambda \in B_k(\rho_k)$, if $\lambda \xrightarrow{\text{sym}, k} \lambda + \alpha$, then $\alpha = \alpha_i$ is some simple root. Moreover, for $\lambda \in B_k(\rho_k)$ we have $\langle \lambda, \alpha_i^\vee \rangle \geq 0$ for any simple root α_i . Hence, for $\lambda \in B_k(\rho_k)$ the symmetric interval-firing relation reduces to

$$\lambda \xrightarrow{\text{sym}, k} \lambda + \alpha_i \text{ for a simple root } \alpha_i, i \in [n] \text{ with } \langle \lambda, \alpha_i^\vee \rangle \leq k - 1.$$

Define $\Psi_k : P \rightarrow P$ by $\Psi_k(\lambda) := -\lambda + \rho_{k+1}$ (so Ψ_k is just a “reflection plus translation”). Then for $\lambda \in \Psi_k^{-1}(B_k(\rho_k)) = B_k(\rho)$ we have

$$\Psi_k(\lambda) \xrightarrow{\text{sym}, k} \Psi_k(\lambda - \alpha_i) \text{ for a simple root } \alpha_i, i \in [n] \text{ with } \langle \Psi_k(\lambda), \alpha_i^\vee \rangle \geq 2.$$

Thus the restriction of $\Psi_k^{-1}(\Gamma_{\text{sym}, k})$ to $B_k(\rho)$ is exactly the same as the restriction of $\Gamma_{\xrightarrow{\mathbf{C}}}$ to $B_k(\rho)$. But every $\lambda \in P$ belongs to $B_k(\rho)$ as $k \rightarrow \infty$. In this way, we can recover Cartan matrix chip-firing as a certain $k \rightarrow \infty$ limit of symmetric interval-firing.

Benkart-Klivans-Reiner [BKR18, Theorem 1.1] showed that the *recurrent configurations* for Cartan matrix chip-firing are $\rho - \omega$ for $\omega \in \Omega_m^0$. Observe $\Psi_k(\rho - \omega) = \eta_k(\omega)$, so these recurrent configurations correspond exactly to the sinks of our full-dimensional and saturated components. In the same way, the 2^n stable configurations in $\mathbb{Z}_{\geq 0}^n$ for Cartan matrix chip-firing correspond to the sinks of our full-dimensional components.

We should stress, however, that confluence is much easier to establish for Cartan matrix chip-firing than for our interval-firing processes: for Cartan matrix chip-firing, confluence holds locally, which ultimately has to do with the fact that simple roots are pairwise non-acute. On the other hand, when firing arbitrary positive roots, confluence need not hold locally because two positive roots may form an acute angle. Hence while Cartan matrix chip-firing describes the limiting behavior of our interval-firing process, it does not explain why the system is confluent from every initial point. Indeed, we could have also obtained Cartan matrix chip-firing by taking the same $k \rightarrow \infty$ limit of the process which has $\lambda \rightarrow \lambda + \alpha$ for $\lambda \in P$, $\alpha \in \Phi^+$ when $\langle \lambda, \alpha^\vee \rangle + 1 \in [-k + 2, k]$, but that process is not confluent.

11. CONFLUENCE OF TRUNCATED INTERVAL-FIRING

So far in this paper we have mostly focused on symmetric interval-firing. We now finally turn to truncated interval-firing. In this section we prove the confluence of $\xrightarrow{\text{tr}, \mathbf{k}}$. Let us start by describing the sinks of $\Gamma_{\text{tr}, \mathbf{k}}$.

Lemma 11.1. *For any $\mathbf{k} \in \mathbb{N}[\Phi]^W$, the sinks of $\Gamma_{\text{tr}, \mathbf{k}}$ are $\{\eta_{\mathbf{k}}(\lambda) : \lambda \in P\}$.*

Proof. Let $\lambda \in P$. Let $\alpha \in \Phi^+$. Note that since $w_\lambda \in W_{I_{\lambda_{\text{dom}}}^0}$, w_λ does not have a descent s_i with $I_{\lambda_{\text{dom}}}^0$ and thus w_λ has no inversions in $\Phi_{I_{\lambda_{\text{dom}}}^0}$. Thus if $\alpha \in w_\lambda(\Phi_{I_{\lambda_{\text{dom}}}^0})$, then $\langle \eta_{\mathbf{k}}(\lambda), \alpha^\vee \rangle = \langle \lambda_{\text{dom}} + \rho_{\mathbf{k}}, w_\lambda^{-1}(\alpha)^\vee \rangle \geq \mathbf{k}(\alpha)$, since $w_\lambda^{-1}(\alpha) \in \Phi^+$. So now consider $\alpha \notin w_\lambda(\Phi_{I_{\lambda_{\text{dom}}}^0})$. Then $w_\lambda^{-1}(\alpha)$ may be positive or negative, but $|\langle \lambda_{\text{dom}}, w_\lambda(\alpha)^\vee \rangle| \geq 1$ (because λ_{dom} has an ω_i coefficient of at least 1 for some $i \notin I_{\lambda_{\text{dom}}}^0$ such that α_i^\vee appears in the expansion of $\pm w_\lambda(\alpha)^\vee$). Hence

$$|\langle \eta_{\mathbf{k}}(\lambda), \alpha^\vee \rangle| = |\langle \lambda_{\text{dom}} + \rho_{\mathbf{k}}, w_\lambda^{-1}(\alpha)^\vee \rangle| \geq \mathbf{k}(\alpha) + 1,$$

which means that $\langle \eta_{\mathbf{k}}(\lambda), \alpha^\vee \rangle \notin [-\mathbf{k}(\alpha), \mathbf{k}(\alpha) - 1]$. So indeed $\eta_{\mathbf{k}}(\lambda)$ is a sink of $\Gamma_{\text{tr}, \mathbf{k}}$.

Now suppose μ is a sink of $\Gamma_{\text{tr}, \mathbf{k}}$. Since $\langle \mu, \alpha^\vee \rangle \notin [-\mathbf{k}(\alpha), \mathbf{k}(\alpha) - 1]$ for all $\alpha \in \Phi^+$, in particular $|\langle \mu, \alpha^\vee \rangle| \geq \mathbf{k}(\alpha)$ for all $\alpha \in \Phi^+$. This means that $\langle \mu_{\text{dom}}, \alpha^\vee \rangle \geq \mathbf{k}(\alpha)$ for all $\alpha \in \Phi^+$. Hence $\mu_{\text{dom}} = \mu' + \rho_{\mathbf{k}}$ for some dominant $\mu' \in P_{\geq 0}$. Suppose to the contrary that w_μ is not the minimal length element of $w_\mu W_{I_{\mu'}^0}$. Then there exists a descent s_i of w_μ with $i \in I_{\mu'}^0$. But then

$$\langle \mu, -w_\mu(\alpha_i)^\vee \rangle = \langle \mu_{\text{dom}}, -\alpha_i^\vee \rangle = -\langle \mu', \alpha_i^\vee \rangle - \langle \rho_{\mathbf{k}}, \alpha_i^\vee \rangle \geq -\mathbf{k}(\alpha_i),$$

and also $\langle \mu, -w_\mu(\alpha_i)^\vee \rangle = -\langle \mu_{\text{dom}}, \alpha_i^\vee \rangle \leq 0$. This would mean μ is not a sink of $\Gamma_{\text{tr}, \mathbf{k}}$, since $-w_\mu(\alpha_i) \in \Phi^+$. So w_μ must be the minimal length element of $w_\mu W_{I_{\mu'}^0}$. This means $w_\mu = w_\lambda$ for some $\lambda \in P$ with $\lambda_{\text{dom}} = \mu'$. And $\mu = w_\mu(\mu_{\text{dom}}) = \lambda + w_\lambda(\rho_{\mathbf{k}}) = \eta_{\mathbf{k}}(\lambda)$, as claimed. \square

We now proceed to prove the confluence of truncated interval-firing. In some sense our proof of confluence here is less satisfactory than the one for symmetric interval-firing because we heavily rely on the diamond lemma, and reduction to rank 2, which is a kind of “trick” that obscures the underlying polytopal geometry (and requires us at one point to use the classification of rank 2 root systems). But we also do crucially use the

permutohedron non-escaping lemma in the following lemma, which says that “small” permutohedra close to the origin are connected components of truncated interval-firing.

Lemma 11.2. *Let $\mathbf{k} \in \mathbb{N}[\Phi]^W$ be good. Then for all $\omega \in \Omega_0^m$, the (translated) discrete permutohedron $\Pi^Q(\rho_{\mathbf{k}}) + \omega$ is a connected component of $\Gamma_{\text{tr}, \mathbf{k}}$ and the unique sink of this connected component is $\rho_{\mathbf{k}} + \omega$.*

Proof. First let us prove a preliminary result: for any $\lambda \in P$ and $\omega \in \Omega_0^m$, we have that $(\lambda - \omega)_{\text{dom}} = \lambda_{\text{dom}} - w(w_\lambda^{-1}(\omega))$ for some $w \in W_{I_{\lambda_{\text{dom}}}^0}$. Indeed, since ω is minuscule or zero, we have that $\langle -w'(\omega), \alpha^\vee \rangle \in \{-1, 0, 1\}$ for any $\alpha \in \Phi$ and any $w' \in W$. Therefore $w_\lambda^{-1}(\lambda - \omega) = \lambda_{\text{dom}} - w_\lambda^{-1}(\omega)$ may not be dominant, but the only α_i for which we have $\langle \lambda_{\text{dom}} - w_\lambda^{-1}(\omega), \alpha_i^\vee \rangle < 0$ must have $i \in I_{\lambda_{\text{dom}}}^0$. Hence, if we let $w \in W_{I_{\lambda_{\text{dom}}}^0}$ be such that $\langle w(w_\lambda^{-1}(\omega)), \alpha_i^\vee \rangle \geq 0$ for all $i \in I_{\lambda_{\text{dom}}}^0$, then $(\lambda - \omega)_{\text{dom}} = \lambda_{\text{dom}} - w(w_\lambda^{-1}(\omega))$ as claimed.

Now let us show that for any $\omega \in \Omega_m^0$, the only sink of $\Gamma_{\text{tr}, \mathbf{k}}$ in $\Pi^Q(\rho_{\mathbf{k}}) + \omega$ is $\rho_{\mathbf{k}} + \omega$. Suppose $\eta_{\mathbf{k}}(\lambda) \in \Pi^Q(\rho_{\mathbf{k}}) + \omega$ for some $\lambda \in P$. This means $\eta_{\mathbf{k}}(\lambda) - \omega \in \Pi^Q(\rho_{\mathbf{k}})$, which means that $(\eta_{\mathbf{k}}(\lambda) - \omega)_{\text{dom}} = \lambda_{\text{dom}} + \rho_{\mathbf{k}} - w(w_\lambda(\omega)) \in \Pi^Q(\rho_{\mathbf{k}})$ for some $w \in W_{I_{\lambda_{\text{dom}}}^0}$ (we are using that $w_{\eta_{\mathbf{k}}(\lambda)} = w_\lambda$, which we have seen before, and that $W_{I_{\lambda_{\text{dom}}}^0} \subseteq W_{I_{\lambda_{\text{dom}} + \rho_{\mathbf{k}}}^0}$). Hence Proposition 2.2 tells us that

$$\rho_{\mathbf{k}} - (\lambda_{\text{dom}} + \rho_{\mathbf{k}} - w(w_\lambda(\omega))) = -(\lambda_{\text{dom}} - \omega) + (w(w_\lambda(\omega)) - \omega) \in Q_{\geq 0}.$$

Now, since $\lambda_{\text{dom}} \in (Q + \omega) \cap P_{\geq 0}$, we know that $\lambda_{\text{dom}} - \omega \in Q_{\geq 0}$ (by one characterization of minuscule weights mentioned in §2). Also, $\omega - w(w_\lambda(\omega)) \in Q_{\geq 0}$ by Proposition 2.2. Hence we conclude that $\lambda_{\text{dom}} = \omega$ and $w(w_\lambda(\omega)) = \omega$. But since we have $w \in W_{I_{\lambda_{\text{dom}}}^0}$, we conclude that $w(w_\lambda(\omega)) = w_\lambda(\omega)$, and thus $w_\lambda(\omega) = \omega$, which forces w_λ to be the identity, i.e., we have $\lambda = \omega$. So indeed the only sink of $\Gamma_{\text{tr}, \mathbf{k}}$ in $\Pi^Q(\rho_{\mathbf{k}}) + \omega$ is $\rho_{\mathbf{k}} + \omega$.

Let us prove the lemma first for $\omega = 0$. Since $\xrightarrow{\text{tr}, \mathbf{k}}$ is terminating by Proposition 4.3, any $\xrightarrow{\text{tr}, \mathbf{k}}$ -firing sequence starting at some $\mu \in \Pi^Q(\rho_{\mathbf{k}})$ has to terminate somewhere. By the permutohedron non-escaping lemma, Lemma 8.2, such a sequence must terminate somewhere inside $\Pi^Q(\rho_{\mathbf{k}})$; and since $\rho_{\mathbf{k}}$ is the only sink in $\Pi^Q(\rho_{\mathbf{k}})$, it must terminate at $\rho_{\mathbf{k}}$. So indeed $\Pi^Q(\rho_{\mathbf{k}})$ is a connected component of $\Gamma_{\text{tr}, \mathbf{k}}$.

Finally, let $\omega \in \Omega_m$ be arbitrary, and let $w \in C$ be the element corresponding to ω under the isomorphism $C \simeq P/Q$. Then by the description of this isomorphism in §5 we get $w(0 - \rho/h) + \rho/h = \omega$, and hence $w(\Pi^Q(\rho_{\mathbf{k}}) - \rho/h) + \rho/h = \Pi^Q(\rho_{\mathbf{k}}) + \omega$. So from the symmetry of $\Gamma_{\text{tr}, \mathbf{k}}^{\text{un}}$ described in Theorem 5.1, we get that $\Pi^Q(\rho_{\mathbf{k}}) + \omega$ is also a connected component of $\Gamma_{\text{tr}, \mathbf{k}}$. \square

Now we consider truncated interval-firing for rank 2 root systems.

Proposition 11.3. *Suppose Φ is of rank 2. Let $\mathbf{k} \in \mathbb{N}[\Phi]^W$. Let $\lambda \in P$ be such that $\langle \lambda, \alpha^\vee \rangle \in [-\mathbf{k}(\alpha), \mathbf{k}(\alpha)]$ and $\langle \lambda, \beta^\vee \rangle \in [-\mathbf{k}(\beta), \mathbf{k}(\beta)]$ for two linearly independent roots $\alpha, \beta \in \Phi$. Suppose that either Φ is simply laced or one of α and β is short and the other is long. Let $\omega \in \Omega_m^0$ be such that $\rho_{\mathbf{k}} - \lambda \in Q + \omega$. Then $\lambda \in \Pi^Q(\rho_{\mathbf{k}}) + \omega$.*

Proof. First let us show $\lambda_{\text{dom}} = c_1\omega_1 + c_2\omega_2$ with $c_1 \in [0, \mathbf{k}(\alpha_1)]$ and $c_2 \in [0, \mathbf{k}(\alpha_2)]$. Observe that $\langle \lambda_{\text{dom}}, w_\lambda(\alpha)^\vee \rangle \in [-\mathbf{k}(\alpha), \mathbf{k}(\alpha)]$ and similarly for β . By replacing α with $-\alpha$ and β with $-\beta$ if necessary, we can assume $\langle \lambda_{\text{dom}}, w_\lambda(\alpha)^\vee \rangle \in [0, \mathbf{k}(\alpha)]$ and similarly for β , and since λ_{dom} is dominant, we are free to assume that $w_\lambda(\alpha)^\vee$ is positive and similarly for β . Note that $w_\lambda(\alpha)^\vee$ and $w_\lambda(\beta)^\vee$ are both nonnegative integer combinations of the simple coroots α_1^\vee and α_2^\vee . Then, since α and β are linearly independent, and since either Φ is simply laced, in which case $\mathbf{k}(\alpha) = \mathbf{k}(\beta) = k$, or one of α, β is short (say e.g. $\mathbf{k}(\alpha) = k_s$) and the other is long (say e.g. $\mathbf{k}(\beta) = k_l$), we can conclude in fact that $\langle \lambda_{\text{dom}}, \alpha_1^\vee \rangle \in [0, \mathbf{k}(\alpha_1)]$ and $\langle \lambda_{\text{dom}}, \alpha_2^\vee \rangle \in [0, \mathbf{k}(\alpha_2)]$.

So indeed, $\lambda_{\text{dom}} = c_1\omega_1 + c_2\omega_2$ with $c_1 \in [0, \mathbf{k}(\alpha_1)]$ and $c_2 \in [0, \mathbf{k}(\alpha_2)]$. If $c_1 = \mathbf{k}(\alpha_1)$ and $c_2 = \mathbf{k}(\alpha_2)$, then $\lambda = \rho_{\mathbf{k}}$ and the proposition is obvious in this case (note that we will have $\omega = 0$). So assume without loss of generality that $c_2 \leq \mathbf{k}(\alpha_2) - 1$.

Let $\lambda' := \lambda - \omega$. We want to show $\lambda' \in \Pi^Q(\rho_{\mathbf{k}})$. As we have seen in the proof of Lemma 11.2, we have $\lambda'_{\text{dom}} = \lambda_{\text{dom}} - w(\omega)$ for some $w \in W$. So let $w \in W$ be such that $\lambda'_{\text{dom}} = \lambda_{\text{dom}} - w(\omega)$ and write $\lambda'_{\text{dom}} = c'_1\omega_1 + c'_2\omega_2$. Since $\langle -w(\omega), \alpha^\vee \rangle \in \{-1, 0, 1\}$ for any $\alpha \in \Phi$, we have $c'_1 \leq \mathbf{k}(\alpha_1) + 1$ and $c'_2 \leq \mathbf{k}(\alpha_2)$. First suppose $c'_1 \leq \mathbf{k}(\alpha_1)$. Together with $c'_2 \leq \mathbf{k}(\alpha_2)$, this implies that $\rho_{\mathbf{k}} - \lambda'_{\text{dom}} \in P_{\geq 0}$, and hence $\rho_{\mathbf{k}} - \lambda'_{\text{dom}} \in Q_{\geq 0}$. Thus we conclude $\lambda' \in \Pi^Q(\rho_{\mathbf{k}})$ by Proposition 2.2.

So suppose that $c'_1 = \mathbf{k}(\alpha_1) + 1$. This means $\langle -w(\omega), \alpha_1^\vee \rangle = 1$. Since ω is the only dominant element of $W(\omega)$, we must actually have that $\langle -w(\omega), \alpha_2^\vee \rangle \leq 0$ and hence $c'_2 \leq \mathbf{k}(\alpha_2) - 1$. Write $\rho_{\mathbf{k}} - \lambda'_{\text{dom}} = a_1\alpha_1 + a_2\alpha_2$ for some integers $a_1, a_2 \in \mathbb{Z}$. Then $c'_1 = \mathbf{k}(\alpha_1) + 1$ and $c'_2 \leq \mathbf{k}(\alpha_2) - 1$ translate to

$$\begin{aligned} 2a_1 + \langle a_2, \alpha_1^\vee \rangle a_2 &= -1; \\ \langle a_1, \alpha_2^\vee \rangle a_1 + 2a_2 &\geq 1. \end{aligned}$$

By the classification of rank 2 root systems we have $\langle a_2, \alpha_1^\vee \rangle, \langle a_1, \alpha_2^\vee \rangle \in \{-1, -2, -3\}$ with at least one of them equal to -1 . It is then not hard to check that all integer solutions $a_1, a_2 \in \mathbb{Z}$ to the above system of inequalities must have $a_1, a_2 \geq 0$. Hence we conclude $\rho_{\mathbf{k}} - \lambda'_{\text{dom}} \in Q_{\geq 0}$, and thus $\lambda' \in \Pi^Q(\rho_{\mathbf{k}})$ by Proposition 2.2. \square

Corollary 11.4. *Suppose Φ is of rank 2. Let $\mathbf{k} \in \mathbb{N}[\Phi]^W$ is good. Then $\xrightarrow{\text{tr}, \mathbf{k}}$ is confluent (and terminating).*

Proof. We know $\xrightarrow{\text{tr}, \mathbf{k}}$ is terminating thanks to Proposition 4.3. Hence by the diamond lemma, Lemma 3.1, it is enough to prove that $\xrightarrow{\text{tr}, \mathbf{k}}$ is locally confluent.

First let us prove this when Φ is simply laced. Suppose $\lambda \xrightarrow{\text{tr}, \mathbf{k}} \lambda + \alpha$ and $\lambda \xrightarrow{\text{tr}, \mathbf{k}} \lambda + \beta$ for $\alpha, \beta \in \Phi^+$. Then by Proposition 11.3 we have that $\lambda \in \Pi^Q(\rho_{\mathbf{k}}) + \omega$ where $\omega \in \Omega_m^0$ is such that $\rho_{\mathbf{k}} - \lambda \in Q + \omega$. But by Lemma 11.1, $\Pi^Q(\rho_{\mathbf{k}}) + \omega$ is a connected component of $\Gamma_{\text{tr}, \mathbf{k}}$ with unique sink $\rho_{\mathbf{k}} + \omega$; since $\xrightarrow{\text{tr}, \mathbf{k}}$ is terminating this means that any $\xrightarrow{\text{tr}, \mathbf{k}}$ -firing sequence starting at λ eventually terminates at $\rho_{\mathbf{k}} + \omega$. Hence we can bring $\lambda + \alpha$ and $\lambda + \beta$ back together again via $\xrightarrow{\text{tr}, \mathbf{k}}$ -firings.

Note that confluence for $\Phi = A_1 \times A_1$ (for any $\mathbf{k} \in \mathbb{N}[\Phi]^W$) reduces to confluence for $\Phi = A_1$, which is trivial. Thus in fact we have proved confluence for all simply laced root systems of rank 2, including those which are not irreducible.

So assume Φ is not simply laced. Suppose $\lambda \xrightarrow{\text{tr}, \mathbf{k}} \lambda + \alpha$ and $\lambda \xrightarrow{\text{tr}, \mathbf{k}} \lambda + \beta$ for $\alpha, \beta \in \Phi^+$. If one of α and β is short and the other is long, then we can apply Proposition 11.3 and Lemma 11.1 as above to conclude that we can bring $\lambda + \alpha$ and $\lambda + \beta$ back together again via $\xrightarrow{\text{tr}, \mathbf{k}}$ -firings. So suppose α and β have the same length. Then let $\tilde{\Phi}$ be the set of all roots in Φ with the same length as α and β . This $\tilde{\Phi}$ will again be a rank 2 root system, and by construction a simply laced one. Hence by the result for simply laced root systems, we know that truncated interval-firing is confluent for $\tilde{\Phi}$; so in particular we can bring $\lambda + \alpha$ and $\lambda + \beta$ back together again via $\xrightarrow{\text{tr}, \mathbf{k}}$ -firings. \square

The confluence of truncated interval-firing for all root systems follows easily from confluence for rank 2 root systems. The following finishes the proof of Theorem 1.1.

Corollary 11.5. *Let $\mathbf{k} \in \mathbb{N}[\Phi]^W$ be good. Then $\xrightarrow{\text{tr}, \mathbf{k}}$ is confluent (and terminating).*

Proof. We know $\xrightarrow{\text{tr}, \mathbf{k}}$ is terminating thanks to Proposition 4.3. Hence by the diamond lemma, Lemma 3.1, it is enough to prove that $\xrightarrow{\text{tr}, \mathbf{k}}$ is locally confluent. Suppose that $\lambda \xrightarrow{\text{tr}, \mathbf{k}} \lambda + \alpha$ and $\lambda \xrightarrow{\text{tr}, \mathbf{k}} \lambda + \beta$ for $\alpha, \beta \in \Phi^+$. Restricting Φ to the span of α and β gives a rank 2 sub-root system, for which we have proved confluence in Corollary 11.4 (as remarked in the proof of that corollary, we in fact proved confluence for *all* rank 2 root systems, including those which are not irreducible). Hence we can bring $\lambda + \alpha$ and $\lambda + \beta$ back together just with truncated interval-firing moves inside that rank 2 sub-root system. \square

Remark 11.6. Our method of proof of confluence for $\xrightarrow{\text{tr}, \mathbf{k}}$ fails when \mathbf{k} is not good; for instance, Lemma 11.2 is not true for general \mathbf{k} , as can be seen in Example 8.3: here $0 \in \Pi^Q(\rho_{\mathbf{k}})$ but 0 does not belong to the connected component of $\Gamma_{\text{tr}, \mathbf{k}}$ containing $\rho_{\mathbf{k}}$. However, we can actually deduce that $\xrightarrow{\text{tr}, \mathbf{k}}$ is confluent for all $\mathbf{k} \in \mathbb{N}[\Phi]^W$ from Corollary 11.5. Indeed, if $\mathbf{k} \in \mathbb{N}[\Phi]^W$ is not good, then $k_s = 0$. But if $k_s = 0$ then we will never be able to fire any short root. In other words, if $k_s = 0$ then truncated interval-firing reduces to truncated interval-firing with respect to the long roots only; and the long roots form a simply laced root system, for which $\xrightarrow{\text{tr}, \mathbf{k}}$ is known to be confluent from Corollary 11.5.

Remark 11.7. It appears that when $\Phi = A_2$ there are no intervals $[a, b]$ for which the relation $\lambda \rightarrow \lambda + \alpha$ for $\lambda \in P$, $\alpha \in \Phi^+$ with $\langle \lambda, \alpha^\vee \rangle + 1 \in [a, b]$ is confluent besides the symmetric and truncated intervals (and this probably would not be too hard to prove). If so, then the same would be true for all irreducible simply laced root systems (except for A_1) because any irreducible root system of rank 3 or greater contains an A_2 sub-root system. This observation also severely restricts possible intervals defining

confluent processes for all root systems, including the non-simply laced ones (although note that central-firing is confluent for $\Phi = B_2$).

Remark 11.8. To any root-firing process \rightarrow on P let us associate the hyperplane arrangement which contains the hyperplane $H = \{v \in V : \langle v, \alpha^\vee \rangle = c\}$ whenever we have a firing move $\lambda \rightarrow \lambda + \alpha$ with $\langle \lambda + \frac{\alpha}{2}, \alpha^\vee \rangle = c$; i.e., we include a hyperplane orthogonal to α at the *midpoint* between λ and $\lambda + \alpha$. As mentioned in the introduction, under this correspondence the symmetric and truncated interval-firing processes correspond to the (extended) Catalan and Shi hyperplane arrangements [PS00, Ath00]. The confluence of symmetric and truncated interval-firing seems like it might have something to do with the freeness of the Catalan and Shi arrangements. *Freeness* is a certain deep algebraic property of hyperplane arrangements introduced by Terao [Ter80]. Freeness of the (extended) Catalan and Shi hyperplane arrangements of a root system was conjectured by Edelman and Reiner [ER96] and proven by Yoshinaga [Yos04] building on work of Terao [Ter02]. Vic Reiner suggested that we look at other free deformations of Coxeter arrangements as a possible source of other confluent root-firing processes. We found one such process which, experimentally, appears confluent: for $\mathbf{k} \in \mathbb{N}[\Phi]$ consider the relation $\lambda \rightarrow \lambda + \alpha$ for $\lambda \in P$, $\alpha \in \Phi^+$ with $\langle \lambda, \alpha^\vee \rangle + 1 \in [-\mathbf{k}(\alpha) + 1, \mathbf{k}(\alpha)]$ if α is long and $\langle \lambda, \alpha^\vee \rangle + 1 \in [-\mathbf{k}(\alpha), \mathbf{k}(\alpha)]$ if α is short. In other words, we use either the truncated or symmetric intervals depending on which Weyl group orbit our root lies in. This process corresponds to a *Shi-Catalan* hyperplane arrangement, as studied by Abe and Terao [AT11]. Other free variants of Coxeter arrangements include the *ideal subarrangements* of Coxeter arrangements [ABC⁺16, AT16], but we have not been able to obtain confluent root-firing processes from these. Note that the freeness of the corresponding hyperplane arrangement certainly does not imply confluence of the root-firing process: for instance, reversing the direction of all the arrows for the truncated interval-firing process yields a process which is not confluent but which corresponds to the same Shi hyperplane arrangement. Nevertheless, it would be very interesting to understand the connection between freeness and confluence further.

Remark 11.9. Under the correspondence between root-firing processes and hyperplane arrangements discussed in Remark 11.8, the central-firing process corresponds not to the central Coxeter arrangement, but rather to the affine *Linial arrangement*. The Linial arrangement has many interesting combinatorial properties (see e.g. [PS00] and [Ath00]), but is not free.

Part 2. Ehrhart-like polynomials

12. EHRHART-LIKE POLYNOMIALS: INTRODUCTION

Continue to fix a root system Φ in vector space V as in the previous part (and retain all the notation from that part). In this part, we investigate the set of weights with given symmetric or truncated interval-firing stabilization. Thus, for good $\mathbf{k} \in \mathbb{N}[\Phi]^W$, we define the *stabilization maps* $s_{\mathbf{k}}^{\text{sym}} : P \rightarrow P$ and $s_{\mathbf{k}}^{\text{tr}} : P \rightarrow P$ by

$$s_{\mathbf{k}}^{\text{sym}}(\mu) = \lambda \Leftrightarrow \text{the } \xrightarrow{\text{sym}, \mathbf{k}} \text{-stabilization of } \mu \text{ is } \eta_{\mathbf{k}}(\lambda);$$

$$s_{\mathbf{k}}^{\text{tr}}(\mu) = \lambda \Leftrightarrow \text{the } \xrightarrow{\text{sym}, \mathbf{k}}\text{-stabilization of } \mu \text{ is } \eta_{\mathbf{k}}(\lambda).$$

These functions are well-defined since the symmetric and truncated interval-firing processes are confluent and terminating (Corollaries 9.2 and 11.5), the stable points of these processes must have the form $\eta_{\mathbf{k}}(\lambda)$ for some $\lambda \in P$ (Lemmas 6.6 and 11.1), and the map $\eta_{\mathbf{k}}$ is injective (Proposition 6.3).

Looking at Example 4.1, one can see that the set $(s_{\mathbf{k}}^{\text{sym}})^{-1}(\lambda)$ (or $(s_{\mathbf{k}}^{\text{tr}})^{-1}(\lambda)$) of weights with interval-firing stabilization $\eta_{\mathbf{k}}(\lambda)$ looks “the same” across all values of \mathbf{k} except that it gets “dilated” as \mathbf{k} is scaled. In analogy with the *Ehrhart polynomial* [Ehr77] of a convex lattice polytope, which counts the number of lattice points in dilations of the polytope, let us define for all $\lambda \in P$ and all good $\mathbf{k} \in \mathbb{N}[\Phi]^W$ the quantities:

$$\begin{aligned} L_{\lambda}^{\text{sym}}(\mathbf{k}) &:= \#(s_{\mathbf{k}}^{\text{sym}})^{-1}(\lambda); \\ L_{\lambda}^{\text{tr}}(\mathbf{k}) &:= \#(s_{\mathbf{k}}^{\text{tr}})^{-1}(\lambda). \end{aligned}$$

Our aim is to show that $L_{\lambda}^{\text{sym}}(\mathbf{k})$ and $L_{\lambda}^{\text{tr}}(\mathbf{k})$ are polynomials in \mathbf{k} . By “polynomial in \mathbf{k} ” we mean that, if Φ is simply laced, then these $L_{\lambda}^{\text{sym}}(\mathbf{k})$ and $L_{\lambda}^{\text{tr}}(\mathbf{k})$ are single-variable polynomials in k , where $\mathbf{k}(\alpha) = k$ for all $\alpha \in \Phi$; and if Φ is non-simply laced, then they are two-variable polynomials in k_s and k_l , where $\mathbf{k}(\alpha) = k_s$ if α is short and $\mathbf{k}(\alpha) = k_l$ if α is long.

We are able to show that the $L_{\lambda}^{\text{sym}}(\mathbf{k})$ are polynomials for all root systems Φ (Theorem 13.3), and we are able to show that the $L_{\lambda}^{\text{tr}}(\mathbf{k})$ are polynomials assuming that Φ is simply laced (Theorem 16.1). In fact, we show that all these polynomials have integer coefficients. Moreover, we conjecture that for all Φ that these $L_{\lambda}^{\text{sym}}(\mathbf{k})$ and $L_{\lambda}^{\text{tr}}(\mathbf{k})$ are polynomials with *nonnegative* integer coefficients.

We refer to these $L_{\lambda}^{\text{sym}}(\mathbf{k})$ and $L_{\lambda}^{\text{tr}}(\mathbf{k})$ as the *symmetric* and *truncated Ehrhart-like polynomials* because they count the size of some discrete subset of lattice points as that set is somehow “dilated.” But it is important to note that the sets $(s_{\mathbf{k}}^{\text{sym}})^{-1}(\lambda)$ and $(s_{\mathbf{k}}^{\text{tr}})^{-1}(\lambda)$ are in general not the set of lattice points of any convex polytope, or indeed, any convex set. This can already be seen in rank 2 (see Example 4.1). Nevertheless, for some special λ (namely, $\lambda \in \Omega_m^0$) the polynomials $L_{\lambda}^{\text{sym}}(\mathbf{k})$ and $L_{\lambda}^{\text{tr}}(\mathbf{k})$ are (essentially) genuine Ehrhart polynomials; and so we do use Ehrhart theory to prove the polynomiality of $L_{\lambda}^{\text{sym}}(\mathbf{k})$ and $L_{\lambda}^{\text{tr}}(\mathbf{k})$. Note that, because they apparently have nonnegative integer coefficients, these polynomials are (as we explain below) most similar to the Ehrhart polynomials of *zonotopes*.

13. SYMMETRIC EHRHART-LIKE POLYNOMIALS

The *Ehrhart polynomial* $L_{\mathcal{P}}(k)$ of a convex lattice polytope \mathcal{P} is a single-variable polynomial in k which satisfies

$$L_{\mathcal{P}}(k) = \text{the number of lattice points in } k\mathcal{P} \text{ (the } k\text{th dilate of } \mathcal{P}\text{)}$$

for all $k \geq 1$. Such polynomials were first investigated by Ehrhart [Ehr77], who proved that they exist for all lattice polytopes. A famous result of Stanley [Sta80, Example 3.1] says that the Ehrhart polynomial of a lattice *zonotope* (i.e., a Minkowski sum of line




FIGURE 10. Paving a polytope plus zonotope.

segments) has nonnegative integer coefficients. A standard way to prove this result is to inductively *pave* the zonotope (see [BR15, §9.2]); this decomposition of a zonotope goes back to Shephard [She74]. In the following theorem we apply this same paving technique to a slightly more general setting: namely, we show that if \mathcal{P} is any fixed convex lattice polytope, and \mathcal{Z} is a lattice zonotope, then for $k \geq 1$ the number of lattice points in $\mathcal{P} + k\mathcal{Z}$ is a polynomial with nonnegative integer coefficients in k . Stanley’s result corresponds to taking \mathcal{P} to be a point. Although the proof is, as mentioned, standard, we have not found this theorem in the Ehrhart theory literature; and it turns out that this result is just what we need to prove that the symmetric Ehrhart-like polynomials $L_\lambda^{\text{sym}}(\mathbf{k})$ exist.

Theorem 13.1. *Let Λ be a lattice in V . Let \mathcal{P} be any convex lattice polytope in V . Let $v_1, \dots, v_m \in \Lambda$ be lattice elements. Then for any $\mathbf{k} = (k_1, \dots, k_m) \in \mathbb{Z}_{\geq 0}^m$ the quantity*

$$\#(\mathcal{P} + k_1[0, v_1] + \dots + k_m[0, v_m]) \cap \Lambda$$

is given by a polynomial in the k_1, \dots, k_m with nonnegative integer coefficients.

Proof. For $X = \{u_1, \dots, u_\ell\} \subseteq V$ linearly independent, a *half-open parallelepiped* with edge set X is a convex set $\mathcal{Z}_X^{h.o.}$ of the form

$$\mathcal{Z}_X^{h.o.} = \sum_{i=1}^{\ell} \begin{cases} [0, u_i] & \text{if } \varepsilon = 1; \\ (0, u_i] & \text{if } \varepsilon = -1, \end{cases}$$

for some choice of sign vector $(\varepsilon_1, \dots, \varepsilon_\ell) \in \{-1, 1\}^\ell$. For $X \subseteq \{v_1, \dots, v_m\}$ let us use $\mathbf{k}X := \{k_i v_i : v_i \in X\}$.

The key idea for this theorem: $\mathcal{P} + k_1[0, v_1] + \dots + k_m[0, v_m]$ can be inductively decomposed (or “paved”) into disjoint pieces that are (up to translation) of the form

$$F + \mathcal{Z}_{\mathbf{k}X}^{h.o.},$$

where $X \subseteq \{v_1, \dots, v_m\}$ is linearly independent and F is an open face of the polytope \mathcal{P} which is affinely independent from $\text{Span}_{\mathbb{R}}(X)$. Figure 10 shows how this is done. Here by “open face” of \mathcal{P} we mean a face minus its relative boundary. Note that vertices have empty relative boundary and hence vertices are open faces. (But observe that Figure 10 is slightly misleading in that we should technically show the whole polytope \mathcal{P} decomposed into its open faces as well; instead the figure shows these pieces grouped

into a single bigger piece.) The proof, by induction on m , that this is possible works in exactly the same way as for paving a zonotope (see [BR15, Lemma 9.1]), so we do not go into the details. Then note that

$$\# \left(\left(F + \mathcal{Z}_{\mathbf{k}X}^{h.o.} \right) \cap \Lambda \right) = \# \left(\left(F + \mathcal{Z}_X^{h.o.} \right) \cap \Lambda \right) \cdot \prod_{v_i \in X} k_i$$

precisely because F is affinely independent from $\mathcal{Z}_{\mathbf{k}X}^{h.o.}$. Hence the desired polynomial in k_1, \dots, k_m indeed exists: it is a sum over the pieces of this decomposition of $\# \left(\left(F + \mathcal{Z}_X^{h.o.} \right) \cap \Lambda \right) \cdot \prod_{v_i \in X} k_i$. (We are implicitly using the fact that this decomposition can be realized in a uniform way across all values of \mathbf{k} .) \square

Corollary 13.2. *For any $\lambda \in P_{\geq 0}$, for all $\mathbf{k} \in \mathbb{N}[\Phi]^W$ the quantity $\#\Pi^Q(\lambda + \rho_{\mathbf{k}})$ is given by a polynomial with nonnegative integer coefficients in \mathbf{k} .*

Proof. We are free to translate $\Pi^Q(\lambda + \rho_{\mathbf{k}})$ so that it contains the origin; i.e., clearly $\#\Pi^Q(\lambda + \rho_{\mathbf{k}})$ is the number of Q -points in $\Pi(\lambda + \rho_{\mathbf{k}}) - \lambda - \rho_{\mathbf{k}}$. One easy consequence of Proposition 2.2 is that $\Pi(\lambda + \mu) = \Pi(\lambda) + \Pi(\mu)$ for dominant weights $\lambda, \mu \in P_{\geq 0}$. Hence, because λ is dominant, we have

$$\Pi(\lambda + \rho_{\mathbf{k}}) - \lambda - \rho_{\mathbf{k}} = (\Pi(\lambda) - \lambda) + (\Pi(\rho_{\mathbf{k}}) - \rho_{\mathbf{k}}).$$

It is well known that the *regular permutohedron* $\Pi(\rho)$ is a zonotope. In Type A, a standard way to prove this fact is to compute the Newton polytope of the Vandermonde determinant in two ways (see [BR15, Theorem 9.4]). The same argument, but with *Weyl's denominator formula* (see [Hum72, §24.3]) in place of the Vandermonde determinant, establishes that $\Pi(\rho) = \sum_{\alpha \in \Phi^+} [-\alpha/2, \alpha/2]$. It is then a simple exercise to show that $\Pi(\rho_{\mathbf{k}}) = \sum_{\alpha \in \Phi^+} \mathbf{k}(\alpha) [-\alpha/2, \alpha/2]$. Hence,

$$\Pi(\lambda + \rho_{\mathbf{k}}) - \lambda - \rho_{\mathbf{k}} = (\Pi(\lambda) - \lambda) + \sum_{\alpha \in \Phi^+} \mathbf{k}(\alpha) [0, -\alpha],$$

and so the desired polynomial indeed exists thanks to Theorem 13.1. \square

We are now ready to prove the first part of Theorem 1.2.

Theorem 13.3. *For any $\lambda \in P$, for good $\mathbf{k} \in \mathbb{N}[\Phi]^W$ the quantity $L_{\lambda}^{\text{sym}}(\mathbf{k})$ is given by a polynomial with integer coefficients in \mathbf{k} .*

Proof. First of all, if λ has $\langle \lambda, \alpha^\vee \rangle = -1$ for some $\alpha \in \Phi^+$ then clearly we can take $L_{\lambda}^{\text{sym}}(\mathbf{k}) := 0$ because, thanks to Lemma 6.6, $\eta_{\mathbf{k}}(\lambda)$ cannot be a sink of $\Gamma_{\text{sym}, \mathbf{k}}$ in this case. So now assume that λ satisfies $\langle \lambda, \alpha^\vee \rangle \neq -1$ for all $\alpha \in \Phi^+$. If $I_{\lambda_{\text{dom}}}^{0,1} \neq [n]$, then, by Theorem 9.1, the connected component of $\Gamma_{\text{sym}, \mathbf{k}}$ containing the sink $\eta_{\mathbf{k}}(\lambda)$ is contained in $w_{\lambda} \Pi_{I_{\lambda_{\text{dom}}}^{0,1}}^Q(\lambda_{\text{dom}})$, which is contained in an affine translate of the strict subspace $\text{Span}_{\mathbb{R}}(w_{\lambda} \Phi_{I_{\lambda_{\text{dom}}}^{0,1}})$. By induction on rank we know the theorem is true for the sub-root system $w_{\lambda} \Phi_{I_{\lambda_{\text{dom}}}^{0,1}}$. Hence, the desired polynomial $L_{\lambda}^{\text{sym}}(\mathbf{k})$ is just the corresponding polynomial for the orthogonal projection of λ onto $\text{Span}_{\mathbb{R}}(w_{\lambda} \Phi_{I_{\lambda_{\text{dom}}}^{0,1}})$. (Here we use the fact that the map $\eta_{\mathbf{k}}$ respects this projection: but this is clear because

the projection of λ and the projection of $w_\lambda(\rho_{\mathbf{k}})$ onto the weight lattice of $w_\lambda\Phi_{I_{\lambda_{\text{dom}}}^{0,1}}$ are both dominant with respect to the choice of $w_\lambda\Phi_{I_{\lambda_{\text{dom}}}^{0,1}}$ as positive roots, which is a subset of Φ^+ by Proposition 6.4.)

So now assume $I_{\lambda_{\text{dom}}}^{0,1} = [n]$. This means that λ is dominant. Let $\mathbf{k} \in \mathbb{N}[\Phi]^W$ be good. Set $S := \{\mu \in P : \langle \mu, \alpha^\vee \rangle \neq -1 \text{ for all } \alpha \in \Phi^+, \eta_{\mathbf{k}}(\mu) \in \Pi^Q(\eta_{\mathbf{k}}(\lambda))\}$; i.e., S is the set of all labels of sinks of $\Gamma_{\text{sym}, \mathbf{k}}$ that are inside of $\Pi^Q(\eta_{\mathbf{k}}(\lambda))$.

We claim that in fact $S = \{\mu \in P : \langle \mu, \alpha^\vee \rangle \neq -1 \text{ for all } \alpha \in \Phi^+, \mu \in \Pi^Q(\lambda)\}$. Indeed, for $\mu \in P$ with $\langle \mu, \alpha^\vee \rangle \neq -1$ for all $\alpha \in \Phi^+$, we have $\eta_{\mathbf{k}}(\mu) \in \Pi^Q(\eta_{\mathbf{k}}(\lambda))$ if and only if $\eta_{\mathbf{k}}(\mu)_{\text{dom}} = \eta_{\mathbf{k}}(\mu_{\text{dom}}) \in \Pi^Q(\eta_{\mathbf{k}}(\lambda))$. By Proposition 2.2, we have that $\eta_{\mathbf{k}}(\mu_{\text{dom}}) \in \Pi^Q(\eta_{\mathbf{k}}(\lambda))$ if and only if $(\lambda + \rho_{\mathbf{k}}) - (\mu_{\text{dom}} + \rho_{\mathbf{k}}) = \lambda - \mu_{\text{dom}} \in Q_{\geq 0}$, which, again by Proposition 2.2, is if and only if $\mu_{\text{dom}} \in \Pi^Q(\lambda)$, that is, if and only if $\mu \in \Pi^Q(\lambda)$. Note that this second description of S is independent of \mathbf{k} . Also note that for all $\mu \neq \lambda \in S$, either $I_{\mu_{\text{dom}}}^{0,1} \neq [n]$ or $\mu = \mu_{\text{dom}}$, and in the latter case we have that μ is strictly less than λ in root order. Now, the permutohedron non-escaping lemma, Lemma 8.2, says that

$$\Pi^Q(\eta_{\mathbf{k}}(\lambda)) = \bigcup_{\mu \in S} (s_{\mathbf{k}}^{\text{sym}})^{-1}(\mu).$$

Hence, rewriting, and taking cardinalities, we get

$$\#(s_{\mathbf{k}}^{\text{sym}})^{-1}(\lambda) = \#\Pi^Q(\eta_{\mathbf{k}}(\lambda)) - \sum_{\mu \neq \lambda \in S} \#(s_{\mathbf{k}}^{\text{sym}})^{-1}(\mu).$$

The quantity $\#\Pi^Q(\eta_{\mathbf{k}}(\lambda))$ is a polynomial in \mathbf{k} with integer coefficients thanks to Theorem 13.2. The quantity $\sum_{\mu \neq \lambda \in S} \#(s_{\mathbf{k}}^{\text{sym}})^{-1}(\mu)$ is a polynomial in \mathbf{k} with integer coefficients by induction on rank and on root order. Since the above equality holds for all good $\mathbf{k} \in \mathbb{N}[\Phi]^W$, we conclude that $L_\lambda^{\text{sym}}(\mathbf{k}) = \#(s_{\mathbf{k}}^{\text{sym}})^{-1}(\lambda)$ is indeed a polynomial in \mathbf{k} with integer coefficients. \square

Table 1 records the polynomials $L_\lambda^{\text{sym}}(\mathbf{k})$ for the irreducible rank 2 root systems, for all $\lambda \in P_{\geq 0}$ with $I_\lambda^{0,1} = [n]$. Compare these polynomials to the graphs of the corresponding symmetric interval-firing processes in Example 4.1.

Remark 13.4. The evaluation of the polynomial $L_\lambda^{\text{sym}}(\mathbf{k})$ for $\mathbf{k} \in \mathbb{N}[\Phi]^W$ not good may not count the number of weights in the connected component of $\Gamma_{\text{sym}, \mathbf{k}}$ containing $\eta_{\mathbf{k}}(\lambda)$. For example, take $\Phi = B_2$ and \mathbf{k} defined by $k_s := 0$ and $k_l := 1$, as in Example 8.3. Then, with $\lambda := 0$, looking at Table 1 we see

$$L_\lambda^{\text{sym}}(\mathbf{k}) = 2k_l^2 + 4k_l k_s + k_s^2 + 2k_l + 2k_s + 1 = 5,$$

while there are only four weights in the connected component of $\Gamma_{\text{sym}, \mathbf{k}}$ containing the sink $\eta_{\mathbf{k}}(\lambda)$. (Here the ‘‘missing’’ weight is of course the origin.)

Conjecture 13.5. *The polynomials $L_\lambda^{\text{sym}}(\mathbf{k})$ have nonnegative integer coefficients.*

When $\lambda \in \Omega_m^0$, we know thanks to Proposition 10.1 that $(s_{\mathbf{k}}^{\text{sym}})^{-1}(\lambda) = \Pi^Q(\lambda + \rho_{\mathbf{k}})$, so Corollary 13.2 implies that Conjecture 13.5 is true in this case. Very recently, the

Φ	λ	$L_\lambda^{\text{sym}}(\mathbf{k})$
A_2	0	$3k^2 + 3k + 1$
A_2	ω_1	$3k^2 + 6k + 3$
A_2	ω_2	$3k^2 + 6k + 3$
A_2	$\omega_1 + \omega_2$	$6k + 6$
B_2	0	$2k_l^2 + 4k_l k_s + k_s^2 + 2k_l + 2k_s + 1$
B_2	ω_1	$4k_l + 4k_s + 4$
B_2	ω_2	$2k_l^2 + 4k_l k_s + k_s^2 + 6k_l + 4k_s + 4$
B_2	$\omega_1 + \omega_2$	$4k_l + 4k_s + 8$
G_2	0	$9k_l^2 + 12k_l k_s + 3k_s^2 + 3k_l + 3k_s + 1$
G_2	ω_1	$12k_l + 6k_s + 6$
G_2	ω_2	$6k_l + 6k_s + 6$
G_2	$\omega_1 + \omega_2$	$6k_l + 6k_s + 12$

TABLE 1. The polynomials $L_\lambda^{\text{sym}}(\mathbf{k})$ for the irreducible rank 2 root systems.

second and fourth authors have proved Conjecture 13.5 in general [HP18]. The first step in their proof of positivity is to give a more refined version of Theorem 13.1 that gives an explicit formula for the number of lattice points in a polytope plus dilating zonotope.

14. CUBICAL SUBCOMPLEXES

In order to proceed further in our investigation of the stabilization maps $s_{\mathbf{k}}^{\text{sym}}$ and $s_{\mathbf{k}}^{\text{tr}}$, and the relation between them, we need to understand a bit more about the connected components of $\Gamma_{\text{sym}, \mathbf{k}}$. We know that the connected component of $\Gamma_{\text{sym}, \mathbf{k}}$ containing the sink $\eta_{\mathbf{k}}(\lambda)$ is contained in the discrete permutohedron $w_\lambda \Pi_{I_{\lambda_{\text{dom}}}}^Q$ (λ_{dom}) (Theorem 9.1); but it can sometimes contain all of this permutohedron (see Proposition 10.1) and can sometimes contain relatively little of it. In this section we will show that there is a small amount of $w_\lambda \Pi_{I_{\lambda_{\text{dom}}}}^Q$ (λ_{dom}) that this connected component must always contain.

The permutohedron $\Pi_I(\lambda)$ has the structure of a polyhedral complex. The *cubical subcomplex* of $\Pi_I(\lambda)$ is the union of all faces of $\Pi_I(\lambda)$ that are cubes; here a *cube* means a product of pairwise orthogonal intervals. We denote the cubical subcomplex by $\square \Pi_I(\lambda)$. Note that every edge is a cube, and hence $\square \Pi_I(\lambda)$ contains at least the 1-skeleton of $\Pi_I(\lambda)$, but it may contain more. We use $\square \Pi_I^Q(\lambda) := \square \Pi_I(\lambda) \cap (Q + \lambda)$.

Proposition 14.1. *Let $\lambda \in P$ with $\langle \lambda, \alpha^\vee \rangle \neq -1$ for all $\alpha \in \Phi^+$ and let $\mathbf{k} \in \mathbb{N}[\Phi]^W$. Let $Y_\lambda := \{\mu \in P : \mu \xleftarrow[\text{sym}, \mathbf{k}]{*} \eta_{\mathbf{k}}(\lambda)\}$ be the connected component of $\Gamma_{\text{sym}, \mathbf{k}}$ containing the sink $\eta_{\mathbf{k}}(\lambda)$. Then Y_λ contains the discrete cubical subcomplex $w_\lambda \square \Pi_{I_{\lambda_{\text{dom}}}}^Q$ ($\eta_{\mathbf{k}}(\lambda_{\text{dom}})$).*

Proof. By the usual projection argument that we have by now carried out many times, we can assume that $I_{\lambda_{\text{dom}}}^{0,1} = [n]$ and consequently that λ is dominant.

For any simple root α_i we have that $\langle \eta_{\mathbf{k}}(\lambda), \alpha_i^\vee \rangle \in \{\mathbf{k}(\alpha), \mathbf{k}(\alpha) + 1\}$. This means that we can “unfire” α_i from $\eta_{\mathbf{k}}(\lambda)$; that is, $\langle \eta_{\mathbf{k}}(\lambda) - \alpha_i, \alpha_i^\vee \rangle \leq \mathbf{k}(\alpha) - 1$, so that there will be an edge $\eta_{\mathbf{k}}(\lambda) - \alpha_i \xrightarrow{\text{sym}, \mathbf{k}} \eta_{\mathbf{k}}(\lambda)$ of $\Gamma_{\text{sym}, \mathbf{k}}$. In fact, we can keep “unfiring” the simple root α_i until we reach $s_{\alpha_i}(\eta_{\mathbf{k}}(\lambda))$; i.e., in $\Gamma_{\text{sym}, \mathbf{k}}$ there are sequence of edges

$$s_{\alpha_i}(\eta_{\mathbf{k}}(\lambda)) \xrightarrow{\text{sym}, \mathbf{k}} s_{\alpha_i}(\eta_{\mathbf{k}}(\lambda)) + \alpha_i \xrightarrow{\text{sym}, \mathbf{k}} \cdots \xrightarrow{\text{sym}, \mathbf{k}} \eta_{\mathbf{k}}(\lambda) - \alpha_i \xrightarrow{\text{sym}, \mathbf{k}} \eta_{\mathbf{k}}(\lambda).$$

(Note that it is possible that $s_{\alpha_i}(\eta_{\mathbf{k}}(\lambda)) = \eta_{\mathbf{k}}(\lambda)$, in which case we would not actually be able to unfire α_i at all). This means that all the $(Q + \eta_{\mathbf{k}}(\lambda))$ -points of the entire edge of $\Pi(\eta_{\mathbf{k}}(\lambda))$ between $\eta_{\mathbf{k}}(\lambda)$ and $s_{\alpha_i}(\eta_{\mathbf{k}}(\lambda))$ are reachable via unfirings from $\eta_{\mathbf{k}}(\lambda)$. Moreover, if α_i and α_j are orthogonal, then unfiring one of these does not affect our ability to unfire the other, and hence in this way we can reach any $(Q + \eta_{\mathbf{k}}(\lambda))$ -point on a face of $\Pi(\eta_{\mathbf{k}}(\lambda))$ that is the orthogonal product of edges coming out of the vertex $\eta_{\mathbf{k}}(\lambda)$ in the direction of a negative simple root. Since in particular $s_i(\eta_{\mathbf{k}}(\lambda))$ is reachable via firings and unfirings from $\eta_{\mathbf{k}}(\lambda)$, by applying the W -symmetry of $\Gamma_{\text{sym}, \mathbf{k}}^{\text{un}}$ (Theorem 5.1) we see that all vertices of $\Pi(\eta_{\mathbf{k}}(\lambda))$ are so reachable. But note that any face of $\Pi(\eta_{\mathbf{k}}(\lambda))$ can be transported via W to a face containing $\eta_{\mathbf{k}}(\lambda)$, such that the edges of this face which contain $\eta_{\mathbf{k}}(\lambda)$ are in the direction of a negative simple root (see the proof of Theorem 7.6). We thus conclude that we can reach any $(Q + \eta_{\mathbf{k}}(\lambda))$ -point on any cubical face of $\Pi(\eta_{\mathbf{k}}(\lambda))$ via firings and unfirings from $\eta_{\mathbf{k}}(\lambda)$. \square

Corollary 14.2. *Let $\lambda \in P$ with $\langle \lambda, \alpha^\vee \rangle \neq -1$ for all $\alpha \in \Phi^+$ and let $\mathbf{k} \in \mathbb{N}[\Phi]^W$. Let $Y_\lambda := \{\mu \in P : \mu \xrightarrow{\text{sym}, \mathbf{k}}^* \eta_{\mathbf{k}}(\lambda)\}$ be the connected component of $\Gamma_{\text{sym}, \mathbf{k}}$ containing the sink $\eta_{\mathbf{k}}(\lambda)$. Then Y_λ contains $w_\lambda W_{I_{\lambda_{\text{dom}}}^{0,1}}(\eta_{\mathbf{k}}(\lambda_{\text{dom}}))$. In the special case $\mathbf{k} = 0$, Y_λ is in fact equal to $w_\lambda W_{I_{\lambda_{\text{dom}}}^{0,1}}(\lambda_{\text{dom}})$.*

Proof. Note that $\square \Pi_I(\mu)$ contains at least the 1-skeleton of $\Pi_I(\mu)$. Thus Y_λ contains $w_\lambda W_{I_{\lambda_{\text{dom}}}^{0,1}}(\eta_{\mathbf{k}}(\lambda_{\text{dom}}))$ by Proposition 14.1. Now suppose $\mathbf{k} = 0$. If $\mu \xrightarrow{\text{sym}, 0} \mu'$ then $\mu' = \mu + \alpha$ for some $\alpha \in \Phi^+$ with $\langle \mu, \alpha^\vee \rangle = -1$, which means that $\mu' = s_\alpha(\mu)$. Hence any two elements in a connected component of $\Gamma_{\text{sym}, 0}$ must be related by a Weyl group element. By Corollary 9.2, each connected component of $\Gamma_{\text{sym}, 0}$ contains only a single sink, and thus the component Y_λ must be exactly $w_\lambda W_{I_{\lambda_{\text{dom}}}^{0,1}}(\lambda_{\text{dom}})$. \square

15. HOW INTERVAL-FIRING COMPONENTS DECOMPOSE

In this section, we study how symmetric and truncated interval-firing components “decompose” into smaller components. Let us explain what we mean by “decompose” more precisely. For any $\mathbf{k} \in \mathbb{N}[\Phi]^W$, $\Gamma_{\text{tr}, \mathbf{k}}$ is a subgraph of $\Gamma_{\text{sym}, \mathbf{k}}$, so the connected components of $\Gamma_{\text{sym}, \mathbf{k}}$ are unions of connected components of $\Gamma_{\text{tr}, \mathbf{k}}$. Similarly, $\Gamma_{\text{sym}, \mathbf{k}}$ is a subgraph of $\Gamma_{\text{tr}, \mathbf{k}+1}$ and so the connected components of $\Gamma_{\text{tr}, \mathbf{k}+1}$ are unions of connected components of $\Gamma_{\text{sym}, \mathbf{k}}$. What we want to show, in both cases, is that the

way these components decompose into smaller components is consistent with the way we label the components by their sinks $\eta_{\mathbf{k}}(\lambda)$.

That the connected components of $\Gamma_{\text{sym},\mathbf{k}}$ break into connected components of $\Gamma_{\text{tr},\mathbf{k}}$ in a way consistent with the map $\eta_{\mathbf{k}}$ turns out to be a simple consequence of the fact that these connected components contain parabolic coset orbits (i.e., a consequence of Corollary 14.2 from the previous section). This is established in the next lemma and corollary.

Lemma 15.1. *For $\lambda, \mu \in P$, if λ and μ belong to the same connected component of $\Gamma_{\text{sym},0}$, then $\eta_{\mathbf{k}}(\lambda)$ and $\eta_{\mathbf{k}}(\mu)$ belong to the same connected component of $\Gamma_{\text{sym},\mathbf{k}}$ for all $\mathbf{k} \in \mathbb{N}[\Phi]^W$.*

Proof. Let $\lambda, \mu \in P$ belong to the same connected component of $\Gamma_{\text{sym},0}$. From Corollary 14.2, we get that $\mu_{\text{dom}} = \lambda_{\text{dom}}$ and also that there is some $w \in w_{\mu}W_{I_{\lambda_{\text{dom}}}^{0,1}}}$ such that $w^{-1}(\lambda)$ is dominant. But by Corollary 6.2 this means $w \in w_{\lambda}W_{I_{\lambda_{\text{dom}}}^0}$, and since the cosets of $W_{I_{\lambda_{\text{dom}}}^{0,1}}$ are unions of cosets of $W_{I_{\lambda_{\text{dom}}}^0}$, this means $w_{\mu}W_{I_{\lambda_{\text{dom}}}^{0,1}} = w_{\lambda}W_{I_{\lambda_{\text{dom}}}^0}$. Thus, Corollary 14.2 tells us that indeed $\eta_{\mathbf{k}}(\lambda)$ and $\eta_{\mathbf{k}}(\mu)$ belong to the same connected component of $\Gamma_{\text{sym},\mathbf{k}}$ for all $\mathbf{k} \in \mathbb{N}[\Phi]^W$. \square

Corollary 15.2. *For all $\mu \in P$ and all good $\mathbf{k} \in \mathbb{N}[\Phi]^W$, we have*

$$s_{\mathbf{k}}^{\text{sym}}(\mu) = s_0^{\text{sym}}(s_{\mathbf{k}}^{\text{tr}}(\mu)).$$

Proof. Since $\Gamma_{\text{tr},\mathbf{k}}$ is a subgraph of $\Gamma_{\text{sym},\mathbf{k}}$, the $\xrightarrow{\text{sym},\mathbf{k}}$ -stabilization of μ is the same as the $\xrightarrow{\text{tr},\mathbf{k}}$ -stabilization of μ . But the $\xrightarrow{\text{tr},\mathbf{k}}$ -stabilization of μ is by definition $\eta_{\mathbf{k}}(\lambda)$ where $\lambda := s_{\mathbf{k}}^{\text{tr}}(\mu)$. Let λ' be the sink of the connected component of $\Gamma_{\text{sym},0}$ containing λ ; hence, $\lambda' = s_0^{\text{sym}}(\lambda)$. Then Lemma 15.1 says that $\eta_{\mathbf{k}}(\lambda')$ is the sink of the connected component of $\Gamma_{\text{sym},\mathbf{k}}$ containing $\eta_{\mathbf{k}}(\lambda)$. In other words, the $\xrightarrow{\text{sym},\mathbf{k}}$ -stabilization of λ is $\eta_{\mathbf{k}}(\lambda')$, i.e., $s_{\mathbf{k}}^{\text{sym}}(\mu) = \lambda' = s_0^{\text{sym}}(s_{\mathbf{k}}^{\text{tr}}(\mu))$. \square

We want an analog of Lemma 15.1 and Corollary 15.2 for truncated interval-firing. But to show that the connected components of $\Gamma_{\text{tr},\mathbf{k}+1}$ break into connected components of $\Gamma_{\text{sym},\mathbf{k}}$ in a way consistent with the map $\eta_{\mathbf{k}}$ turns out to be much more involved. In fact, for technical reasons, we are able to achieve this only assuming that Φ is simply laced. Nevertheless, the first few steps towards giving truncated analogs of Lemma 15.1 and Corollary 15.2 do not require the assumption that Φ be simply laced, so we state them for general Φ .

Proposition 15.3. *Let $\lambda \in P$ be such that $\langle \lambda, \alpha^{\vee} \rangle \neq -1$ for all $\alpha \in \Phi^+$. Suppose that $\lambda \xrightarrow{\text{tr},1} \lambda + \beta$ for some $\beta \in \Phi^+$. Then $\lambda \xrightarrow{\text{tr},1} \lambda + w_{\lambda}(\alpha_i)$ for some simple root α_i . Moreover, in this case we have $\eta_{\mathbf{k}}(\lambda) \xrightarrow{\text{tr},\mathbf{k}+1} \eta_{\mathbf{k}}(\lambda) + w_{\lambda}(\alpha_i)$ for all $\mathbf{k} \in \mathbb{N}[\Phi]^W$.*

Proof. If $\langle \lambda, \alpha^{\vee} \rangle \neq -1$ for all $\alpha \in \Phi^+$, but $\lambda \xrightarrow{\text{tr},1} \lambda + \beta$ for some $\beta \in \Phi^+$, this must mean that $\langle \lambda, \beta^{\vee} \rangle = 0$. Applying w_{λ}^{-1} , we get $\langle w_{\lambda}^{-1}(\lambda), w_{\lambda}^{-1}(\beta)^{\vee} \rangle = 0$. Since $w_{\lambda}^{-1}(\beta)^{\vee}$ is either a positive sum or a negative sum of simple coroots, and because $w_{\lambda}^{-1}(\lambda) = \lambda_{\text{dom}}$

is dominant, this means there is some simple root α_i such that $\langle w_\lambda^{-1}(\lambda), \alpha_i^\vee \rangle = 0$. But then $\langle \lambda, w_\lambda(\alpha_i)^\vee \rangle = 0$. And note by Proposition 6.4 that indeed $w_\lambda(\alpha)$ is positive.

To prove the last sentence of the proposition: note that

$$\langle \eta_{\mathbf{k}}(\lambda), w_\lambda(\alpha_i)^\vee \rangle = \langle \lambda + w_\lambda(\rho_{\mathbf{k}}), w_\lambda(\alpha_i)^\vee \rangle = \langle w_\lambda^{-1}(\lambda), \alpha_i^\vee \rangle + \langle \rho_{\mathbf{k}}, \alpha_i^\vee \rangle = 0 + \mathbf{k}(\alpha) = \mathbf{k}(\alpha);$$

so indeed, $\eta_{\mathbf{k}}(\lambda) \xrightarrow{\text{tr}, \mathbf{k}+1} \eta_{\mathbf{k}}(\lambda) + w_\lambda(\alpha_i)$. \square

Proposition 15.4. *Let $\lambda \in P$ be a weight such that $\langle \lambda, \alpha^\vee \rangle \neq -1$ for all $\alpha \in \Phi^+$. Let $\mathbf{k} \in \mathbb{N}[\Phi]^W$ be good, with $\mathbf{k} \geq 1$. Let $\mu \in w_\lambda \Pi_{I_{\lambda_{\text{dom}}}^{0,1}}^Q(\eta_{\mathbf{k}}(\lambda_{\text{dom}}))$. Then μ and $\eta_{\mathbf{k}}(\lambda)$ belong to the same connected component of $\Gamma_{\text{tr}, \mathbf{k}+1}$.*

Proof. First let us prove this proposition when λ is dominant and $I_\lambda^0 = [n]$. In this case, $\eta_{\mathbf{k}}(\lambda) \in \Pi(\rho_{\mathbf{k}+1})$. Let $\omega \in \Omega_m^0$ be such that $\lambda \in Q + \rho + \omega$. Note that, since $\mathbf{k} \geq 1$, $\eta_{\mathbf{k}}(\lambda) - \omega$ is still dominant; hence, because $P_{\geq 0}^{\mathbb{R}} \subseteq Q_{\geq 0}^{\mathbb{R}}$, we get that $\eta_{\mathbf{k}}(\lambda) - \omega \in \Pi(\rho_{\mathbf{k}+1})$ by Proposition 2.2. But then by definition of ω we have that $\eta_{\mathbf{k}}(\lambda) \in \Pi^Q(\rho_{\mathbf{k}+1}) + \omega$. Thus by Lemma 11.2 the connected component of $\Gamma_{\text{tr}, \mathbf{k}+1}$ that $\eta_{\mathbf{k}}(\lambda)$ belongs to is $\Pi^Q(\rho_{\mathbf{k}+1}) + \omega$. By Corollary 14.2, the connected component of $\Gamma_{\text{sym}, \mathbf{k}}$ that $\eta_{\mathbf{k}}(\lambda)$ belongs to contains the Weyl orbit $W(\eta_{\mathbf{k}}(\lambda))$. Hence also the the connected component of $\Gamma_{\text{tr}, \mathbf{k}+1}$ that $\eta_{\mathbf{k}}(\lambda)$ belongs to contains $W(\eta_{\mathbf{k}}(\lambda))$. But this connected component is, as mentioned, $\Pi^Q(\rho_{\mathbf{k}+1}) + \omega$; in particular, it is a convex set intersected with $Q + \eta_{\mathbf{k}}(\lambda)$. Since μ belongs to the convex hull of $W(\eta_{\mathbf{k}}(\lambda))$ and belongs to the coset $Q + \eta_{\mathbf{k}}(\lambda)$, this means that $\mu \in \Pi^Q(\rho_{\mathbf{k}+1}) + \omega$. So indeed μ and $\eta_{\mathbf{k}}(\lambda)$ belong to the same connected component of $\Gamma_{\text{tr}, \mathbf{k}+1}$ in this case.

Now let us address general λ . Note that $w_\lambda \Phi_{I_{\lambda_{\text{dom}}}^{0,1}}^+$ is a choice of positive roots for the sub-root system $w_\lambda \Phi_{I_{\lambda_{\text{dom}}}^{0,1}}$. Moreover, by Proposition 6.4, $w_\lambda \Phi_{I_{\lambda_{\text{dom}}}^{0,1}}^+$ is a subset of positive roots. Hence any truncated interval-firing move (with parameter $\mathbf{k} + 1$) we can carry out in $w_\lambda \Phi_{I_{\lambda_{\text{dom}}}^{0,1}}$ with choice of positive roots $w_\lambda \Phi_{I_{\lambda_{\text{dom}}}^{0,1}}^+$, we can actually carry out in the original root system Φ . But then note that $\langle \lambda, w_\lambda(\alpha_i)^\vee \rangle \in \{0, 1\}$ for all $i \in I_{\lambda_{\text{dom}}}^{0,1}$; hence the result follows from the previous paragraph by orthogonally projecting λ and μ onto $\text{Span}_{\mathbb{R}}(w_\lambda \Phi_{I_{\lambda_{\text{dom}}}^{0,1}})$. \square

The strategy will be to use Proposition 15.3 to say that whenever we have a $\xrightarrow{\text{tr}, 1}$ -move from a sink of $\Gamma_{\text{sym}, 0}$, we have a corresponding $\xrightarrow{\text{tr}, \mathbf{k}+1}$ -move from the corresponding sink of $\Gamma_{\text{sym}, \mathbf{k}}$; then we will apply Proposition 15.4 to say that that move actually gets us “trapped” in the correct connected component of $\Gamma_{\text{tr}, \mathbf{k}+1}$. But we have reached the point where to carry out this strategy we must assume that Φ is simply laced.

Proposition 15.5. *Suppose that Φ is simply laced. Let $\mu \in P_{\geq 0}$ be dominant. Suppose $\mu \xrightarrow{\text{tr}, 1} \lambda$ where $\lambda = \mu + \alpha_i$ for a simple root α_i . Then $\lambda \in W_{I_{\lambda_{\text{dom}}}^{0,1}}(\lambda_{\text{dom}})$.*

Proof. If μ is dominant but $\mu \xrightarrow{\text{tr}, 1} \lambda$, this must mean that $\langle \mu, \alpha_i^\vee \rangle = 0$. Let Φ' be the irreducible sub-root system of $\Phi_{I_\mu^0}$ that contains α_i . Let θ' be the highest root of Φ' .

We claim that $\lambda_{\text{dom}} = \mu + \theta'$. First of all, because Φ' is also simply laced, the Weyl group W' of Φ' acts transitively on Φ' so that there is some $w \in W'$ with $w(\theta') = \alpha_i$. But $W' \subseteq W_{I_\mu^0}$, the stabilizer of μ , so we indeed have $w(\mu + \theta') = \mu + \alpha_i = \lambda$. Now, why is $\mu + \theta'$ dominant? Let D be the Dynkin diagram of Φ (which is just an undirected graph since Φ is simply laced). For $I \subseteq [n]$ use $D[I]$ to denote the restriction of the Dynkin diagram to the vertices in I . Note that $\Phi' = \Phi_I$ where I is (the set of vertices of) the connected component of $D[I_\mu^0]$ containing α_i . Hence $\theta' = \sum_{j \in I} c_j \alpha_j$ for some coefficients c_j . First of all, θ' is dominant in Φ' , so if $j \in I$ then $\langle \theta', \alpha_j^\vee \rangle \geq 0$ and hence certainly $\langle \mu + \theta', \alpha_j^\vee \rangle \geq 0$. Now suppose $j \notin I$ and j is not adjacent in D to any vertex in I ; then clearly $\langle \theta', \alpha_j^\vee \rangle = 0$ and so again $\langle \mu + \theta', \alpha_j^\vee \rangle \geq 0$. Finally, suppose $j \notin I$ but j is adjacent in D to some vertex in I ; then, since Φ is simply laced and θ' is a positive root of Φ , we certainly have $\langle \theta', \alpha_j^\vee \rangle \geq -1$; but $\langle \lambda, \alpha_j^\vee \rangle \geq 1$ since $j \notin I_\mu^0$, and thus $\langle \lambda + \theta', \alpha_j^\vee \rangle \geq 0$. So indeed $\mu + \theta'$ is dominant and so $\lambda_{\text{dom}} = \mu + \theta'$, as claimed.

Now, suppose for a moment that $\Phi' \neq A_1$. Then, writing $\theta' = \sum_{j=1}^n c_j \omega_j$, we will have that $c_j \in \{0, 1\}$ for all $j \in I$; this can be seen for instance by noting that these coefficients c_j are precisely the number of edges between j and the ‘‘affine node’’ in the affine Dynkin diagram extending $D[I]$ (see [Bou02, VI, §3]). This means that we have $W' \subseteq W_{I_{\lambda_{\text{dom}}}}^{0,1}$, and so $w(\lambda_{\text{dom}}) = \lambda$ for some $w \in W_{I_{\lambda_{\text{dom}}}}^{0,1}$; or in other words, we have $\lambda \in W_{I_{\lambda_{\text{dom}}}}^{0,1}(\lambda_{\text{dom}})$. On the other hand, if $\Phi' = A_1$, then actually $\theta' = \alpha_i$ and so $\lambda = \lambda_{\text{dom}}$ and the claim is clear. \square

Remark 15.6. Note that Proposition 15.5 is in general false when Φ is not simply laced. For example, take $\Phi = B_2$. Then, with $\mu := 0$ and $\lambda := \alpha_1$ (the long simple root, with numbering as in Figure 2), we have $\mu \xrightarrow{\text{tr},1} \lambda$ but $\lambda \notin W_{I_{\lambda_{\text{dom}}}}^{0,1}(\lambda_{\text{dom}})$.

Proposition 15.7. *Suppose that Φ is simply laced. Let $\mu \in P_{\geq 0}$ be dominant. Suppose that $\mu \xrightarrow{\text{tr},1} \lambda$ where $\lambda = \mu + \alpha_i$ for a simple root α_i . Then $\eta_k(\mu) + \alpha \in \Pi_{I_{\lambda_{\text{dom}}}}^{Q,0,1}(\eta_k(\lambda_{\text{dom}}))$ for all $k \geq 0$.*

Proof. The statement in the case $k = 0$ follows immediately from Proposition 15.5; so assume $k \geq 1$. By Proposition 15.5 we have that $\lambda \in W_{I_{\lambda_{\text{dom}}}}^{0,1}(\lambda_{\text{dom}})$, which means, by Proposition 2.2, that $\lambda_{\text{dom}} - \lambda$ is a nonnegative sum of simple roots in $I_{\lambda_{\text{dom}}}^{0,1}$. Since μ is dominant we have $\eta_k(\mu) = \mu + k\rho$. Then note that $\eta_k(\mu) + \alpha = \lambda + k\rho = \mu + k\rho + \alpha$ is actually dominant as well, because μ is dominant, and $k\rho + \alpha$ is dominant since Φ is simply laced. Further, observe that $\eta_k(\lambda_{\text{dom}}) - (\eta_k(\mu) + \alpha) = \lambda_{\text{dom}} - \lambda$. But then the fact that $\eta_k(\lambda_{\text{dom}}) - (\eta_k(\mu) + \alpha)$ is a nonnegative sum of simple roots in $I_{\lambda_{\text{dom}}}^{0,1}$, together with the fact that $\eta_k(\mu) + \alpha$ is dominant, implies, via Proposition 2.2, that we have $\eta_k(\mu) + \alpha \in \Pi_{I_{\lambda_{\text{dom}}}}^{Q,0,1}(\eta_k(\lambda_{\text{dom}}))$. \square

Proposition 15.8. *Suppose that Φ is simply laced. Let $\mu \in P$ satisfy $\langle \mu, \alpha^\vee \rangle \neq -1$ for all $\alpha \in \Phi^+$. Suppose that $\mu \xrightarrow{\text{tr},1} \lambda$ where $\lambda = \mu + w_\mu(\alpha_i)$ for some simple root α_i . Then for all $k \geq 0$, $\eta_k(\mu)$ and $\eta_k(\lambda)$ belong to the same connected component of $\Gamma_{\text{tr},k+1}$.*

Proof. If $k = 0$ the claim is obvious. So assume $k \geq 1$.

Let λ' be the sink of the connected component of $\Gamma_{\text{sym},0}$ containing λ ; hence by Corollary 14.2, we have that $\lambda' \in w_\lambda W_{I_{\lambda_{\text{dom}}}^{0,1}}(\lambda_{\text{dom}})$, so in particular $\lambda'_{\text{dom}} = \lambda_{\text{dom}}$. Now, if $\langle \mu, \alpha^\vee \rangle \neq -1$ for all $\alpha \in \Phi^+$ and $\mu \xrightarrow{\text{tr},1} \lambda$ this means that $\langle \mu, w_\mu(\alpha_i)^\vee \rangle = 0$. Hence we also have $\mu_{\text{dom}} \xrightarrow{\text{tr},1} \mu_{\text{dom}} + \alpha_i$. Then $\mu_{\text{dom}} + \alpha_i \in W_{I_{\lambda_{\text{dom}}}^{0,1}}(\lambda_{\text{dom}})$ by Proposition 15.5; and so by applying w_μ we get $\lambda \in w_\mu W_{I_{\lambda_{\text{dom}}}^{0,1}}(\lambda_{\text{dom}})$. This implies $\lambda' \in w_\mu W_{I_{\lambda_{\text{dom}}}^{0,1}}(\lambda_{\text{dom}})$, so that $(w_\mu w)^{-1}(\lambda')$ is dominant for some $w \in W_{I_{\lambda_{\text{dom}}}^{0,1}}$. But because of Corollary 6.2 that means that $w_\mu w = w_{\lambda'} w'$ for some $w' \in W_{I_{\lambda_{\text{dom}}}^0}$.

By Proposition 15.7 we get that $\eta_k(\mu_{\text{dom}}) + \alpha_i \in \Pi_{I_{\lambda_{\text{dom}}}^{0,1}}^Q(\eta_k(\lambda_{\text{dom}}))$. By applying w_μ we get $\eta_k(\mu) + w_\mu(\alpha_i) \in w_\mu \Pi_{I_{\lambda_{\text{dom}}}^{0,1}}^Q(\eta_k(\lambda_{\text{dom}}))$. Note that since $w \in W_{I_{\lambda_{\text{dom}}}^{0,1}}$, we have that $w_\mu \Pi_{I_{\lambda_{\text{dom}}}^{0,1}}^Q(\eta_k(\lambda_{\text{dom}})) = w_\mu w \Pi_{I_{\lambda_{\text{dom}}}^{0,1}}^Q(\eta_k(\lambda_{\text{dom}}))$. Similarly, $w' \in W_{I_{\lambda_{\text{dom}}}^0} \subseteq W_{I_{\lambda_{\text{dom}}}^{0,1}}$ implies that $w_{\lambda'} w' \Pi_{I_{\lambda_{\text{dom}}}^{0,1}}^Q(\eta_k(\lambda_{\text{dom}})) = w_{\lambda'} \Pi_{I_{\lambda_{\text{dom}}}^{0,1}}^Q(\eta_k(\lambda_{\text{dom}}))$. Hence, we can conclude that $\eta_k(\mu) + w_\mu(\alpha_i) \in w_{\lambda'} \Pi_{I_{\lambda_{\text{dom}}}^{0,1}}^Q(\eta_k(\lambda_{\text{dom}}))$. Since λ' is a sink of $\Gamma_{\text{sym},0}$ (and thus, by Lemma 6.6, satisfies $\langle \lambda', \alpha^\vee \rangle \neq -1$ for all $\alpha \in \Phi^+$), we can apply Proposition 15.4 to conclude that $\eta_k(\lambda')$ and $\eta_k(\mu) + w_\mu(\alpha_i)$ belong to the same connected component of $\Gamma_{\text{tr},k+1}$.

But since λ and λ' belong to the same connected component of $\Gamma_{\text{sym},0}$, Lemma 15.1 tells us that $\eta_k(\lambda)$ and $\eta_k(\lambda')$ belong to the same connected component of $\Gamma_{\text{sym},k}$, and hence also belong to the same connected component of $\Gamma_{\text{tr},k+1}$. Then note by Proposition 15.3 that we have $\eta_k(\mu) \xrightarrow{\text{tr},k+1} \eta_k(\mu) + w_\mu(\alpha_i)$, so $\eta_k(\mu)$ and $\eta_k(\mu) + w_\mu(\alpha_i)$ belong to the same connected component of $\Gamma_{\text{tr},k+1}$. Putting it all together, $\eta_k(\mu)$ and $\eta_k(\lambda)$ belong to the same connected component of $\Gamma_{\text{tr},k+1}$, as claimed. \square

Finally, we are able to prove the desired analogs of Lemma 15.1 and Corollary 15.2 in the simply laced case.

Lemma 15.9. *Suppose that Φ is simply laced. For $\lambda, \mu \in P$, if λ and μ belong to the same connected component of $\Gamma_{\text{tr},1}$, then $\eta_k(\lambda)$ and $\eta_k(\mu)$ belong to the same connected component of $\Gamma_{\text{tr},k+1}$ for all $k \geq 0$.*

Proof. Clearly it suffices to prove this when λ is a sink of $\Gamma^{\text{tr},1}$. So let us describe one way to compute the $\xrightarrow{\text{tr},k+1}$ -stabilization of $\eta_k(\mu)$. If μ is not a sink of $\Gamma_{\text{sym},0}$, then by Lemma 15.1 we know that $\eta_k(\mu)$ is in the same connected component of $\Gamma_{\text{sym},k}$ as $\eta_k(\mu')$, where μ' is the sink of the component of $\Gamma_{\text{sym},0}$ containing μ ; so then to compute the $\xrightarrow{\text{tr},k+1}$ -stabilization of $\eta_k(\mu)$ we instead compute the $\xrightarrow{\text{tr},k+1}$ -stabilization of $\eta_k(\mu')$. So now assume that μ is a sink of $\Gamma_{\text{sym},0}$. Then, if μ is not a sink of $\Gamma_{\text{tr},1}$, by Proposition 15.3 there is a simple root α_i with $\mu \xrightarrow{\text{tr},1} \mu'$ where $\mu' = \mu + w_\mu(\alpha_i)$. By Proposition 15.8 we get that $\eta_k(\mu)$ and $\eta_k(\mu')$ are in the same connected component

of $\Gamma_{\text{tr},k+1}$; so again to compute the $\xrightarrow{\text{tr},k+1}$ -stabilization of $\eta_k(\mu)$ we instead compute the $\xrightarrow{\text{tr},k+1}$ -stabilization of $\eta_k(\mu')$. Because $\xrightarrow{\text{tr},1}$ is terminating, this procedure will eventually terminate; in fact, it must terminate at computing the $\xrightarrow{\text{tr},k+1}$ -stabilization of $\eta_k(\mu)$ where μ is a sink of $\Gamma_{\text{tr},1}$. But there is only one sink of the connected component of $\Gamma_{\text{tr},1}$ containing μ , namely, λ ; so the lemma is proved. \square

Corollary 15.10. *Suppose that Φ is simply laced. Then for all $\mu \in P$ and all $k \geq 0$, we have*

$$s_{k+1}^{\text{tr}}(\mu) = s_1^{\text{tr}}(s_k^{\text{sym}}(\mu)).$$

Proof. This follows from Lemma 15.9 in the same way that Corollary 15.2 follows from Lemma 15.1. Since $\Gamma_{\text{sym},k}$ is a subgraph of $\Gamma_{\text{sym},k+1}$, the $\xrightarrow{\text{tr},k+1}$ -stabilization of μ is the same as the $\xrightarrow{\text{tr},k+1}$ -stabilization of the $\xrightarrow{\text{sym},k}$ -stabilization of μ . But the $\xrightarrow{\text{sym},k}$ -stabilization of μ is by definition $\eta_k(\lambda)$ where $\lambda := s_k^{\text{sym}}(\mu)$. Let $\eta_1(\lambda')$ be the sink of the connected component of $\Gamma_{\text{tr},1}$ containing λ ; hence, $\lambda' = s_1^{\text{tr}}(\lambda)$. Then Lemma 15.1 says that $\eta_k(\eta_1(\lambda')) = \eta_{k+1}(\lambda')$ (this equality follows from Proposition 6.3) is the sink of the connected component of $\Gamma_{\text{tr},k+1}$ containing $\eta_k(\lambda)$. In other words, the $\xrightarrow{\text{tr},k+1}$ -stabilization of λ is $\eta_{k+1}(\lambda')$, i.e., $s_{k+1}^{\text{tr}}(\mu) = \lambda' = s_1^{\text{tr}}(s_k^{\text{sym}}(\mu))$. \square

We expect that (with the appropriate care regarding the goodness of $\mathbf{k} \in \mathbb{N}[\Phi]^W$) Lemma 15.1 and Corollary 15.2 should hold in the non-simply laced case as well, but, as we mentioned in Remark 15.6, our method of proof does not work there.

16. TRUNCATED EHRHART-LIKE POLYNOMIALS

The existence of the truncated Ehrhart-like polynomials, in the simply laced case, follows easily from the fact that truncated components decompose into symmetric ones in a consistent way (together with the existence of the symmetric Ehrhart-like polynomials).

Theorem 16.1. *Suppose that Φ is simply laced. Then, for any $\lambda \in P$, for all $k \geq 1$ the quantity $L_\lambda^{\text{tr}}(k)$ is given by a polynomial in k with integer coefficients.*

Proof. By Corollary 15.10, for any $k \geq 1$ and any $\lambda \in P$ we have

$$\begin{aligned} \#(s_k^{\text{tr}})^{-1}(\lambda) &= \#(s_{k-1}^{\text{sym}})^{-1}((s_1^{\text{tr}})^{-1}(\lambda)) \\ &= \sum_{\mu \in (s_1^{\text{tr}})^{-1}(\lambda)} L_\mu^{\text{sym}}(k-1). \end{aligned}$$

The right-hand side of this expression is an evaluation of a polynomial (with integer coefficients) because of Theorem 13.3. Since this identity holds for all $k \geq 1$, we conclude that the desired polynomial $L_\lambda^{\text{tr}}(k)$ does exist. \square

This finishes the proof of Theorem 1.2.

Conjecture 16.2. *For any Φ and $\lambda \in P$, for all good $\mathbf{k} \in \mathbb{N}[\Phi]^W$ the quantity $L_\lambda^{\text{tr}}(\mathbf{k})$ is given by a polynomial with nonnegative integer coefficients in \mathbf{k} .*

λ	$L_\lambda^{\text{tr}}(k)$
0	$3k^2 + 3k + 1$
ω_1	$3k^2 + 3k + 1$
$-\omega_1 + \omega_2$	$2k + 1$
$-\omega_2$	$k + 1$
ω_2	$3k^2 + 3k + 1$
$\omega_1 - \omega_2$	$2k + 1$
$-\omega_1$	$k + 1$
$\omega_1 + \omega_2$	$2k + 1$
$-\omega_1 + 2\omega_2$	$k + 1$
$2\omega_1 - \omega_2$	$k + 1$
$-2\omega_1 + \omega_2$	$k + 1$
$\omega_1 - 2\omega_2$	$k + 1$
$-\omega_1 - \omega_2$	1

 TABLE 2. The polynomials $L_\lambda^{\text{tr}}(k)$ for $\Phi = A_2$.

Note that the fact we can take $\mathbf{k} = 0$ in Conjecture 16.2 means that the constant term of the $L_\lambda^{\text{tr}}(\mathbf{k})$ polynomials should be 1 (which, compared to the symmetric polynomials, makes them even more like Ehrhart polynomials of zonotopes). Strictly speaking, our Theorem 16.1 does not establish that these polynomials have constant term 1 even in the simply laced case.

Remark 16.3. Table 2 records the polynomials $L_\lambda^{\text{tr}}(k)$ for $\Phi = A_2$, for all $\lambda \in P$ with $I_{\lambda_{\text{dom}}}^{0,1} = [n]$. Compare these polynomials to the graphs of the A_2 truncated interval-firing processes in Example 4.1. In agreement with Conjecture 16.2, all these polynomials have constant coefficient 1. Note that, for $\lambda \in P$ with $L_\lambda^{\text{sym}}(k) \neq 0$, the constant term of $L_\lambda^{\text{sym}}(k)$ is by definition equal to the number of vertices in the connected component of $\Gamma_{\text{sym},0}$ containing λ , which by Lemma 15.1 is also equal to the number of connected components of $\Gamma_{\text{tr},k}$ contained in the connected component of $\Gamma_{\text{sym},k}$ with sink $\eta_k(\lambda)$ for all $k \geq 0$.

We know that Conjecture 16.2 holds for $\lambda \in \Omega_m^0$. That is because, for $\lambda \in \Omega_m^0$, Lemma 11.2 tells us that $(s_{\mathbf{k}}^{\text{tr}})^{-1}(\lambda) = \Pi(\rho_{\mathbf{k}}) + \lambda$, and hence $\#(s_{\mathbf{k}}^{\text{tr}})^{-1}(\lambda)$ is literally the Ehrhart polynomial of a zonotope.

Polynomials with nonnegative integer coefficients occupy a special place in algebraic combinatorics. Of course it would be great, in the course of positively resolving Conjectures 13.5 and 16.2, to also give a combinatorial interpretation of the coefficients of the coefficients of these polynomials. (In fact, this is precisely what is done in [HP18].) It would also be extremely interesting to relate these polynomials to the representation theory or algebraic geometry attached to the root system Φ , and establish positivity in that way. These polynomials arose for us in the course of a purely combinatorial investigation, but it is hard to imagine that they do not have some deeper significance if they indeed have nonnegative integer coefficients.

Remark 16.4. It is also worth considering how the stabilization maps $s_{\mathbf{k}}^{\text{sym}}$ and $s_{\mathbf{k}}^{\text{tr}}$ interact with the symmetries of $\Gamma_{\text{sym},\mathbf{k}}^{\text{un}}$ and $\Gamma_{\text{tr},\mathbf{k}}^{\text{un}}$ coming from Theorem 5.1. For the symmetric stabilization maps: if $\lambda \in P$ and $w \in W^{\lambda_{\text{dom}}^{0,1}}$, then it is not hard to deduce from Lemma 15.1 that

$$(s_{\mathbf{k}}^{\text{sym}})^{-1}(w(\lambda)) = w((s_{\mathbf{k}}^{\text{sym}})^{-1}(\lambda))$$

for all good $\mathbf{k} \in \mathbb{N}[\Phi]^W$. Of course this implies that

$$L_{w(\lambda)}^{\text{sym}}(\mathbf{k}) = L_{\lambda}^{\text{sym}}(\mathbf{k}),$$

in this case. Meanwhile, it appears that if $w \in C \subseteq W$ and $\varphi: P \rightarrow P$ is the affine map $\varphi: v \mapsto w(v - \rho/h) + \rho/h$, then

$$(s_{\mathbf{k}}^{\text{tr}})^{-1}(\varphi(\lambda)) = \varphi((s_{\mathbf{k}}^{\text{tr}})^{-1}(\lambda))$$

for all $\lambda \in P$ and all good $\mathbf{k} \in \mathbb{N}[\Phi]^W$. But even in the simply laced case, where we have Lemma 15.9 at our disposal, in order to conclude that $s_{\mathbf{k}}^{\text{tr}}$ indeed respects the symmetry φ in this way, we would need to know that this is the case for $\mathbf{k} = 1$; and, as we mention in the next section, we do not currently have a great understanding of $\Gamma_{\text{tr},1}$. So to show that the truncated stabilization maps and polynomials have the expected symmetries coming from the subgroup C would require some more work.

17. ITERATIVE DESCRIPTIONS OF THE STABILIZATION

Finally, let us focus a little more on what our decomposition results tell us about the relationship between the polynomials $L_{\lambda}^{\text{sym}}(\mathbf{k})$ and $L_{\lambda}^{\text{tr}}(\mathbf{k})$, and between the stabilization map $s_{\mathbf{k}}^{\text{sym}}$ and $s_{\mathbf{k}}^{\text{tr}}$. So, let us assume that Φ is simply laced for the remainder of this section. It is clear that Corollaries 15.2 and 15.10 imply the following identities relating these polynomials for all $\lambda \in P$ and all $k \geq 1$:

$$\begin{aligned} L_{\lambda}^{\text{sym}}(k) &= \sum_{\mu \in (s_0^{\text{sym}})^{-1}(\lambda)} L_{\mu}^{\text{tr}}(k); \\ L_{\lambda}^{\text{tr}}(k) &= \sum_{\mu \in (s_1^{\text{tr}})^{-1}(\lambda)} L_{\mu}^{\text{sym}}(k-1). \end{aligned}$$

What is more, these corollaries also immediately imply some striking, iterative descriptions of the stabilization functions:

Corollary 17.1. *Suppose that Φ is simply laced. Then for all $\mu \in P$ and all $k \geq 1$:*

- $s_1^{\text{sym}}(\mu) = s_0^{\text{sym}}(s_1^{\text{tr}}(\mu));$
- $s_k^{\text{sym}}(\mu) = (s_1^{\text{sym}})^k(\mu);$
- $s_k^{\text{tr}}(\mu) = s_1^{\text{tr}}((s_1^{\text{sym}})^{k-1}(\mu)).$

Corollary 17.1 says that the information of all of the stabilization maps is contained just in s_0^{sym} and s_1^{tr} . Now, s_0^{sym} is pretty simple to understand: for example, its fibers are just parabolic Weyl coset orbits (see Corollary 14.2). So somehow all of the complexity of all truncated and symmetric interval-firing processes (or, at least all the complexity related to *stabilization* for these interval-firing processes) is contained just in $\Gamma_{\text{tr},1}$.




FIGURE 11. The map $s_1^{\text{sym}}: P \rightarrow P$ for $\Phi = A_1$. We write $L_\lambda^{\text{sym}}(k)$ above each weight $\lambda \in P$.

Admittedly, we do not understand $\Gamma_{\text{tr},1}$ very well. It would be very interesting, for example, to try to find an explicit description of the connected components of $\Gamma_{\text{tr},1}$.

Finally, we end the paper by discussing another surprising consequence of Corollary 17.1: for all $\lambda \in P$ and all $k \geq 1$,

$$\#((s_1^{\text{sym}})^k)^{-1}(\lambda) = L_\lambda^{\text{sym}}(k).$$

In other words, we have a map $f: X \rightarrow X$ from some discrete set to itself, such that the sizes $\#(f^k)^{-1}(x)$ of fibers of iterates of this map are given by polynomials (in k) for every point $x \in X$. In fact, we have many such maps, one for each simply laced root system. This is a very special property for a self-map of a discrete set to have. In the next two examples we show what this looks like in the simplest cases.

Example 17.2. Although we have so far been eschewing one-dimensional examples, in fact s_1^{sym} is interesting even for A_1 . Figure 11 depicts s_1^{sym} for $\Phi = A_1$. Of course in this picture we draw an arrow from μ to λ to mean that $s_1^{\text{sym}}(\mu) = \lambda$. The colors of the vertices correspond to classes of weights modulo the root lattice. We write the polynomials $L_\lambda^{\text{sym}}(k)$ above the weights in this figure. One can verify by hand that in this case $\#((s_1^{\text{sym}})^k)^{-1}(\lambda) = L_\lambda^{\text{sym}}(k)$ for all $\lambda \in P$ and all $k \geq 0$.

Example 17.3. Note that when $\Phi = A_2$, we have $\rho \in Q$ and hence s_1^{sym} preserves the root lattice and so descends to a map $s_1^{\text{sym}}: Q \rightarrow Q$. Figure 12 depicts $s_1^{\text{sym}}: Q \rightarrow Q$ for $\Phi = A_2$. (As with our previous drawings for rank 2 interval-firing processes, we of course only depict the “interesting,” finite portion of this function near the origin.) Compare this figure to the symmetric interval-firing graphs for A_2 in Example 4.1 and the polynomials $L_\lambda^{\text{sym}}(k)$ for A_2 recorded in Table 1. Observe that indeed $((s_1^{\text{sym}})^k)^{-1}(0) = \Pi^Q(k\rho)$ for all $k \geq 1$. Also observe that $((s_1^{\text{sym}})^k)^{-1}(\alpha_1 + \alpha_2)$ is the set of Q -lattice points on the boundary of $\Pi((k+1)\rho)$.

REFERENCES

- [ABC⁺16] Takuro Abe, Mohamed Barakat, Michael Cuntz, Torsten Hoge, and Hiroaki Terao. The freeness of ideal subarrangements of Weyl arrangements. *J. Eur. Math. Soc. (JEMS)*, 18(6):1339–1348, 2016.
- [ALS⁺89] Richard Anderson, László Lovász, Peter Shor, Joel Spencer, Éva Tardos, and Shmuel Winograd. Disks, balls, and walls: analysis of a combinatorial game. *Amer. Math. Monthly*, 96(6):481–493, 1989.
- [AT11] Takuro Abe and Hiroaki Terao. The freeness of Shi-Catalan arrangements. *European J. Combin.*, 32(8):1191–1198, 2011.
- [AT16] Takuro Abe and Hiroaki Terao. Free filtrations of affine Weyl arrangements and the ideal-Shi arrangements. *J. Algebraic Combin.*, 43(1):33–44, 2016.




FIGURE 12. The map $s_1^{\text{sym}}: Q \rightarrow Q$ for $\Phi = A_2$. The origin is the central point (i.e., the one with a loop). The root $\alpha_1 + \alpha_2$ is the point immediately north-east of the origin.

- [Ath00] Christos A. Athanasiadis. Deformations of Coxeter hyperplane arrangements and their characteristic polynomials. In *Arrangements—Tokyo 1998*, volume 27 of *Adv. Stud. Pure Math.*, pages 1–26. Kinokuniya, Tokyo, 2000.
- [BB05] Anders Björner and Francesco Brenti. *Combinatorics of Coxeter groups*, volume 231 of *Graduate Texts in Mathematics*. Springer, New York, 2005.
- [Big99] N. L. Biggs. Chip-firing and the critical group of a graph. *J. Algebraic Combin.*, 9(1):25–45, 1999.
- [BKR18] Georgia Benkart, Caroline Klivans, and Victor Reiner. Chip firing on Dynkin diagrams and McKay quivers. *Math. Z.*, 2018. <https://doi.org/10.1007/s00209-017-2034-5>.
- [BL92] Anders Björner and László Lovász. Chip-firing games on directed graphs. *J. Algebraic Combin.*, 1(4):305–328, 1992.
- [BLS91] Anders Björner, László Lovász, and Peter W. Shor. Chip-firing games on graphs. *European J. Combin.*, 12(4):283–291, 1991.
- [BN07] Matthew Baker and Serguei Norine. Riemann-Roch and Abel-Jacobi theory on a finite graph. *Adv. Math.*, 215(2):766–788, 2007.
- [Bou02] Nicolas Bourbaki. *Lie groups and Lie algebras. Chapters 4–6*. Elements of Mathematics (Berlin). Springer-Verlag, Berlin, 2002. Translated from the 1968 French original by Andrew Pressley.
- [BR15] Matthias Beck and Sinai Robins. *Computing the continuous discretely*. Undergraduate Texts in Mathematics. Springer, New York, second edition, 2015. Integer-point enumeration in polyhedra, With illustrations by David Austin.
- [BTW87] Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality: An explanation of the $1/f$ noise. *Phys. Rev. Lett.*, 59:381–384, 1987.

- [CP18] Scott Corry and David Perkinson. Divisors and sandpiles. In the *AMS Student Mathematical Library* series, 2018. Book in progress; draft available online at http://people.reed.edu/~davidp/divisors_and_sandpiles/.
- [Dha90] Deepak Dhar. Self-organized critical state of sandpile automaton models. *Phys. Rev. Lett.*, 64(14):1613–1616, 1990.
- [Dha99] Deepak Dhar. The abelian sandpile and related models. *Physica A: Statistical Mechanics and its Applications*, 263(1):4 – 25, 1999.
- [Ehr77] E. Ehrhart. *Polynômes arithmétiques et méthode des polyèdres en combinatoire*. Birkhäuser Verlag, Basel-Stuttgart, 1977. International Series of Numerical Mathematics, Vol. 35.
- [Eng75] Arthur Engel. The probabilistic abacus. *Educational Studies in Mathematics*, 6(1):1–22, Mar 1975.
- [Eng76] Arthur Engel. Why does the probabilistic abacus work? *Educational Studies in Mathematics*, 7(1):59–69, Jul 1976.
- [ER96] P. H. Edelman and V. Reiner. Free arrangements and rhombic tilings. *Discrete Comput. Geom.*, 15(3):307–340, 1996.
- [ES03] Pavel Etingof and Elisabetta Strickland. Lectures on quasi-invariants of Coxeter groups and the Cherednik algebra. *Enseign. Math. (2)*, 49(1-2):35–65, 2003.
- [FL16] Matthew Farrell and Lionel Levine. CoEulerian graphs. *Proc. Amer. Math. Soc.*, 144(7):2847–2860, 2016.
- [Gab93] Andrei Gabrielov. Asymmetric abelian avalanches and sandpile. Preprint, *Mathematical Sciences Institute*, Cornell University; available online at <https://www.math.purdue.edu/~agabriel/asym.pdf>, 1993.
- [GHMP17] Pavel Galashin, Sam Hopkins, Thomas McConville, and Alexander Postnikov. Root system chip-firing II: central-firing. Eprint published online at [arXiv:1708.04849](https://arxiv.org/abs/1708.04849), 2017.
- [GK08] Andreas Gathmann and Michael Kerber. A Riemann-Roch theorem in tropical geometry. *Math. Z.*, 259(1):217–230, 2008.
- [GK15] Johnny Guzmán and Caroline Klivans. Chip-firing and energy minimization on M-matrices. *J. Combin. Theory Ser. A*, 132:14–31, 2015.
- [HMP17] Sam Hopkins, Thomas McConville, and James Propp. Sorting via chip-firing. *Electron. J. Comb.*, 24(3):P3.13, 2017.
- [HP18] Sam Hopkins and Alexander Postnikov. A positive formula for the Ehrhart-like polynomials from root system chip-firing. Eprint published online at [arXiv:1803.08472](https://arxiv.org/abs/1803.08472), 2018.
- [Hue80] Gérard Huet. Confluent reductions: Abstract properties and applications to term rewriting systems. *J. ACM*, 27(4):797–821, October 1980.
- [Hum72] James E. Humphreys. *Introduction to Lie algebras and representation theory*. Springer-Verlag, New York-Berlin, 1972. Graduate Texts in Mathematics, Vol. 9.
- [KT15] Viktor Kiss and Lilla Tóthmérés. Chip-firing games on Eulerian digraphs and **NP**-hardness of computing the rank of a divisor on a graph. *Discrete Appl. Math.*, 193:48–56, 2015.
- [LP10] Lionel Levine and James Propp. What is ... a sandpile? *Notices Amer. Math. Soc.*, 57(8):976–979, 2010.
- [LP12] Thomas Lam and Alexander Postnikov. Alcoved polytopes II. Eprint published online at [arXiv:1202.4015](https://arxiv.org/abs/1202.4015), 2012.
- [LPS16] Lionel Levine, Wesley Pegden, and Charles K. Smart. Apollonian structure in the Abelian sandpile. *Geom. Funct. Anal.*, 26(1):306–336, 2016.
- [MZ08] Grigory Mikhalkin and Ilia Zharkov. Tropical curves, their Jacobians and theta functions. In *Curves and abelian varieties*, volume 465 of *Contemp. Math.*, pages 203–230. Amer. Math. Soc., Providence, RI, 2008.
- [New42] M. H. A. Newman. On theories with a combinatorial definition of “equivalence.”. *Ann. of Math. (2)*, 43:223–243, 1942.
- [PS00] Alexander Postnikov and Richard P. Stanley. Deformations of Coxeter hyperplane arrangements. *J. Combin. Theory Ser. A*, 91(1-2):544–597, 2000. In memory of Gian-Carlo Rota.

- [PS04] Alexander Postnikov and Boris Shapiro. Trees, parking functions, syzygies, and deformations of monomial ideals. *Trans. Amer. Math. Soc.*, 356(8):3109–3142, 2004.
- [PS13] Wesley Pegden and Charles K. Smart. Convergence of the Abelian sandpile. *Duke Math. J.*, 162(4):627–642, 2013.
- [She74] G. C. Shephard. Combinatorial properties of associated zonotopes. *Canad. J. Math.*, 26:302–321, 1974.
- [Spe86] J. Spencer. Balancing vectors in the max norm. *Combinatorica*, 6(1):55–65, 1986.
- [Sta80] Richard P. Stanley. Decompositions of rational convex polytopes. *Ann. Discrete Math.*, 6:333–342, 1980. Combinatorial mathematics, optimal designs and their applications (Proc. Sympos. Combin. Math. and Optimal Design, Colorado State Univ., Fort Collins, Colo., 1978).
- [Ste98] John R. Stembridge. The partial order of dominant weights. *Adv. Math.*, 136(2):340–364, 1998.
- [Ter80] Hiroaki Terao. Arrangements of hyperplanes and their freeness. I and II. *J. Fac. Sci. Univ. Tokyo Sect. IA Math.*, 27(2):293–312, 1980.
- [Ter02] Hiroaki Terao. Multiderivations of Coxeter arrangements. *Invent. Math.*, 148(3):659–674, 2002.
- [Ver75] Daya-Nand Verma. The rôle of affine Weyl groups in the representation theory of algebraic Chevalley groups and their Lie algebras. In *Lie groups and their representations (Proc. Summer School, Bolyai János Math. Soc., Budapest, 1971)*, pages 653–705. Halsted, New York, 1975.
- [Yos04] Masahiko Yoshinaga. Characterization of a free arrangement and conjecture of Edelman and Reiner. *Invent. Math.*, 157(2):449–454, 2004.

E-mail address: galashin@mit.edu

E-mail address: shopkins@mit.edu

E-mail address: thomasmc@mit.edu

E-mail address: apost@math.mit.edu

DEPARTMENT OF MATHEMATICS, MASSACHUSETTS INSTITUTE OF TECHNOLOGY, CAMBRIDGE, MA 02139, USA