MOVE AND CONFIGURATION POSETS

CAROLINE KLIVANS, PATRICK LISCIO

Abstract. Lower Locally Distributive (LLD) lattices are a class of lattices that describe many types of combinatorial processes. While LLD lattices are traditionally defined in terms of their join-irreducibles, we provide a new characterization of LLD lattices in terms saturated L-colorings that we call S-colorings. S- and L-colorings color the edges of an LLD lattice's Hasse diagram in a way that has a natural interpretation in terms of moves of a locally confluent process. We show that each color of an S-coloring corresponds to the "j^{th} move of a given type" in a process, allowing us to prove results about locally confluent processes by instead proving them in the context of LLD lattices. We use LLD lattice-based methods to prove new results on root system central-firing, as well as providing novel proofs of key results in labeled chip-firing and domino tilings.

1. Introduction

The goal of this paper is to examine a class of nondeterministic combinatorial processes that can be represented using a poset P. Such a process consists of **configurations**, represented by elements in P, and **moves**, represented by edges in the Hasse diagram of P. We usually consider processes with a fixed initial configuration, represented by a maximum element in P. Final configurations correspond to minimal elements in P. A single instance of the process is represented by a downward path through the Hasse diagram of P.

Conversely, we introduce the idea that any finite poset P with a maximum element can be represented as a process. Elements in P correspond to configurations, and edges in the Hasse diagram of P correspond to moves. Again, any downward path through the Hasse diagram corresponds to one possible instance of the nondeterministic process.

As one might expect, certain properties of a given poset P translate into properties of the corresponding process. We are particularly concerned with properties related to **confluence**, in which a process involving many possible sequences of moves, still has a fixed final configuration for every initial configuration. If a poset corresponding to a confluent process has a unique maximum element, then it also has a unique minimum element.

Some processes have a related property called **local confluence**. If two moves are available from the same configuration in a locally confluent process, then they can be performed in either order, producing the same resulting configuration. It is known that any locally confluent, terminating process with a unique initial configuration must also have a unique final configuration. Furthermore, all sequences of moves from the initial configuration to the final one must be of the same length. The result is often referred to as the Diamond Lemma due to its implications on the corresponding poset of configurations. The idea of local confluence is demonstrated in Figure 1.

The elements a, b, c, and d in Figure 1 form a diamond in the Hasse diagram of P. In many processes, the two paths through the diamond from a to d come from performing the same two moves in two different orders. We say that a poset generated by a locally confluent process has the **diamond property**.

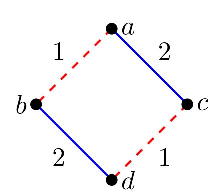

Figure 1. A poset illustrating the local confluence property

This paper deals with a specific class of posets with the diamond property, known as lower locally distributive (LLD) lattices, as well as the behavior of processes corresponding to LLD lattices. Processes exhibiting LLD lattice structure include chip-firing [14] [7], root system central-firing from [8] (shown here), and the sand pile model from [2] (see [12]). All distributive lattices are LLD, so processes involving graph orientations, bipartite matchings, domino tilings, and graph spanning trees [16] also have LLD configuration posets.

We use properties of LLD lattices to prove results on several locally confluent processes. In particular, we simplify the proof of Klivans and Liscio [11] of a theorem of Hopkins, McConville, and Propp [9] that labeled chip-firing beginning with an even number of chips must result in a sorted final configuration. These methods can also be used to prove sorting in labeled chip-firing on more general classes of initial configurations [12]. We extend these methods to prove conjectures of Galashin et al. [8] on sorting in type A, B, and D root system central-firing.

We then turn to domino tilings, where we provide a new proof of a result of Saldanha et al. [17], see also Thurston [20], that any shortest path between flip-connected tilings must consist of the same flip moves, although possibly in a different order.

Theorem 7.5. *All shortest paths between two domino tilings on a simply connected region consist of the same flip moves in possibly different orders.*

We prove results on the above processes by considering the poset of configurations for each process, and then coloring the edges of the resulting Hasse diagrams according to certain rules. These colorings are based on the idea that locally confluent diamonds represent two moves occurring in two different orders. Felsner and Knauer [7] formalize this idea on LLD lattices, by defining a coloring, called an L-coloring, that is compatible with this diamond property.

The first contribution of the current paper is to provide an equivalent characterization of LLD lattices that uses a more refined coloring to more directly address the relationship between LLD lattices and their underlying processes. A saturated L-coloring, or S-coloring, assigns the same color to two edges in the Hasse diagram if the edges appear across from each other in some diamond in the Hasse diagram of P. Subject to that constraint, the S-coloring uses as many colors as possible, ensuring that any instance of the process contains exactly one edge of each color. Our first main result concerns the existence and uniqueness of S-colorings.

Theorem 4.11. *Each LLD lattice P has a unique S-coloring up to isomorphism.*

An S-coloring is of particular interest because of its relationship with the join-irreducibles of the lattice. In particular, each color in an S-coloring contains a single edge that has a join-irreducible as its upper endpoint. See figure 2. One consequence of Theorem 4.11 is that all downward paths between any elements u and v must contain the same colors, and any complete downward path must contain exactly one edge of each color. When viewing a poset as representing a process, we can view the colors of the S-coloring as defining the types of moves that must occur throughout the process.

This allows us to define a partial ordering on the colors in an S-coloring, which we call the **color poset**. Given an S-coloring c with colors c^1 and c^2, we say that $c^1 \geq c^2$ if every complete downward path contains its edge of color c^1 before its edge of color c^2. Our next result relates the color poset to the poset of join-irreducibles:

Theorem 4.12. *The color poset of an LLD lattice is isomorphic to its poset of join irreducibles.*

In many processes, the colors of an S-coloring have a more natural interpretation: each color corresponds to "the j^{th} move of a certain type." Thus, the color poset represents a sort of "move poset," which dictates the orders in which certain types of moves may occur in a process. We provide a general way to prove correspondences between colors and elements of move posets, by providing a more precise relationship between colors of S-colorings and L-colorings.

Corollary 4.17. *Every color in an S-coloring corresponds to a move in the underlying process of the form "j^{th} move of type t," where j is a fixed integer and t corresponds to a color in an L-coloring.*

Since a color in an L-coloring often corresponds to a certain type t of move in a process, the refined colors of the S-coloring allow us to refer to "the n colors of type t." Corollary 4.17 extends this idea further,

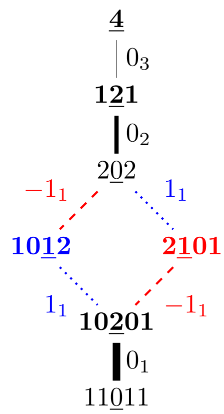

FIGURE 2. The poset of configurations for unlabeled chip-firing on the line beginning with $n = 4$ chips at the origin. The number of chips at the origin is underlined, and the join-irreducible configurations are in bold. There is an L-coloring with 3 colors in which each color corresponds to which site is fired. Here, red (dashed) represents firing moves at -1, black (solid) represents firing moves at 0, and blue (dashed) represents firing moves at 1. The L-coloring can is subdivided into an S-coloring, in which each color corresponds to a specific move k_j. The thick black edge corresponds to move 0_1, the medium black edge corresponds to move 0_2, and the thin black edge corresponds to move 0_3.

showing that the colors of the S-coloring represent "the $1^{st}, 2^{nd}, \ldots n^{th}$ moves of type t" rather than just "the n moves of type t." Thus, an S-coloring provides a well-defined notion of the ordering of the moves in the underlying process.

In Section 4.2, we prove the following result about shortest paths through LLD lattices and the colors that appear along those paths.

Theorem 4.21. *In an S-colored LLD lattice, all shortest paths between two fixed elements consist of edges of the same colors.*

Theorem 4.21 is used to prove the Theorem 7.5 on domino tilings: shortest paths between flip-connected domino tilings, using as few flip moves as possible, must involve the same flip moves in a possibly different order.

S-colorings also allow us to compare LLD lattices to other types of lattices that are also defined by edge labelings, such as supersolvable lattices, with their corresponding S_n EL-labelings.

In section 3, we provide preliminaries on LLD lattices and L-colorings, as well as define an S-coloring. In section 4, we prove the main results on S-colorings. We show that S-colorings are precisely saturated L-colorings, and that each color corresponds to a single join-irreducible of the configuration poset. We introduce the color poset and show that it is isomorphic to the poset of join-irreducibles. Finally, we show that all shortest paths between two configurations must consist of edges of the same colors.

In section 5, we apply our results on S-colorings and LLD lattices to chip-firing and use them to prove sorting results on labeled chip-firing. In section 6, we apply our results to resolve open problems in root system central-firing. In section 7, we interpret the S-colorings of distributive lattices generated by graph orientations, bipartite matchings, and spanning trees. Section 7 also includes our result on shortest paths between domino tilings. In section 8, we relate LLD lattices and S-colorings to supersolvable lattices and their corresponding S_n EL-labelings.

2. ACKNOWLEDGEMENTS

The second author was supported by the National Science Foundation under the RTG Grant DMS-2038039.

3. Preliminaries

Consider a poset P with Hasse diagram $H = (V, E)$, where V is the set of elements of P, or equivalently, the vertices of the Hasse diagram, and E is the set of edges of the Hasse diagram of P, such that $(u, v) \in E$ if $u \gtrdot v$ in P. Define E' to be the inverse of E, so that $(v, u) \in E'$ if $(u, v) \in E$. If $u \gtrdot v$, then there is a **downward edge** from u to v in E, and there is an **upward edge** from v to u in E'.

Define a **downward path** through the Hasse diagram H as a sequence downward edges $(v_0, v_1), (v_1, v_2), \dots (v_{k-1}, v_k)$. If P has a maximum element $\mathbf{1}$ and a minimum element $\mathbf{0}$, define a **complete downward path** to be a downward path in which $v_0 = \mathbf{1}$ and $v_k = \mathbf{0}$. Note that for elements u and v in P, $u \geq v$ if and only if there is a downward path from u to v in H.

We are concerned primarily with a class of posets known as **lower locally distributive (LLD) lattices**. LLD lattices are characterized by two equivalent definitions.

Definition 3.1. A poset P is a **lower locally distributive (LLD) lattice** if either of the following equivalent conditions is met:

(1) P is a lattice, and each element $u \in P$ has a unique minimal representation as a join of join-irreducibles.
(2) For each element $u \in P$, the interval between u and the meet of its lower covers is a boolean lattice.

The inverse of an LLD lattice is called upper locally distributive (ULD). ULD lattices can similarly be defined in terms of meet-irreducibles, or in terms of the join of upper covers of each element. A poset is a distributive lattice if and only if it is LLD and ULD.

More can be said about the representations in definition (1). Define J_u as the set of all join-irreducibles less than or equal to the element u. If $u \gtrdot v$, then $J_v \subseteq J_u$, and $|J_u \backslash J_v| = 1$. As a consequence, every LLD lattice is ranked, with height equal to the number of join-irreducibles. Furthermore, each downward edge (u, v) in E can be associated with a single join-irreducible j such that $\{j\} = J_u \backslash J_v$. Define the **join change** of a downward edge (u, v) to be the single join-irreducible in $J_u \backslash J_v$. As $\mathbf{1}$ is the join of all join-irreducibles, and $\mathbf{0}$ is an empty join of join-irreducibles, every complete downward path through H must contain one edge with each join change.

Define a **coloring** c of a poset P with Hasse diagram $H = (V, E)$ to be a function $c : E \cup E' \to [n]$ such that $c\big((v, u)\big) = c\big((u, v)\big)$ for all $(u, v) \in E$. c induces a partition of $E \cup E'$ into n subsets of the form $c^i = \{e \in E \cup E' : c(e) = i\}$. Each of these n subsets is a **color** ("c^i is a color from the coloring c"). Thus, coloring c has n colors.

Given two colorings c and c', c' is a **refinement** of c if each color in c' is a subset of a color in c. c' is a **proper refinement** of c if c' is a refinement of c and c' has more nonempty colors than c. This paper deals with two colorings of LLD lattices.

Definition 3.2. [7] Given a poset P, an ***L*-coloring** of P is a coloring c such that:

(1) For all $u, v, w \in P$ with $(u, v) \in E$, $(u, w) \in E$, it must be the case that $c\big((u, v)\big) \neq c\big((u, w)\big)$.
(2) For all $u, v, w \in P$ with $(u, v) \in E$, $(u, w) \in E$, there exists a $z \in P$ such that $(v, z) \in E$, $(w, z) \in E$, $c\big((u, v)\big) = c\big((w, z)\big)$, and $c\big((u, w)\big) = c\big((v, z)\big)$.

A poset P that admits an L-coloring is known as an ***L*-poset**. An example of an LLD Lattice with an L-coloring is shown in Figure 3.

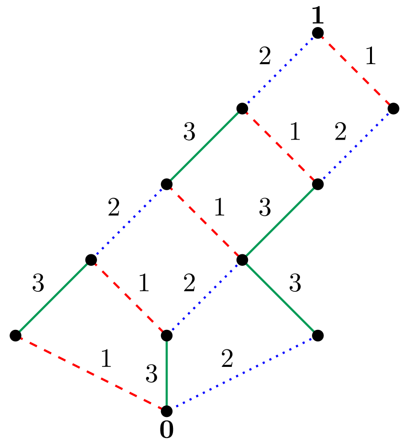

Figure 3. The Hasse diagram of an LLD Lattice P, along with an L-coloring. Each downward path from $\mathbf{1}$ to $\mathbf{0}$ consists of 1 red (dashed) edge, 2 blue (dotted) edges, and 2 green (solid) edges.

Theorem 3.3 ([7], Theorem 1)**.** *A poset is an LLD lattice if and only if it is an L-poset with a global maximum.*

We introduce a closely related coloring.

Definition 3.4. An **S-coloring** of P is a coloring c such that

(1) For any u, v, w, z such that $(u, v), (w, z) \in E$ and $v \geq w$, $c\big((u, v)\big) \neq c\big((w, z)\big)$.
(2) For all $u, v, w \in P$ with $(u, v), (u, w) \in E$, there exists a $z \in P$ such that $(v, z) \in E$, $(w, z) \in E$, $c\big((u, v)\big) = c\big((w, z)\big)$, and $c\big((u, w)\big) = c\big((v, z)\big)$.

A poset is **S-colorable** if it admits an S-coloring. We refer to condition (1) as the "downward path condition," as it states that any downward path through the Hasse diagram cannot contain multiple edges from the same color. Condition (2) of the S-coloring definition is precisely the same as the corresponding L-coloring condition. We will use S-colorings to provide a characterization of LLD lattices, by showing that a poset is an LLD lattice if and only if it admits an S-coloring and has a global maximum (see Lemma 4.8). An S-coloring of the lattice from Figure 3 is shown in Figure 4:

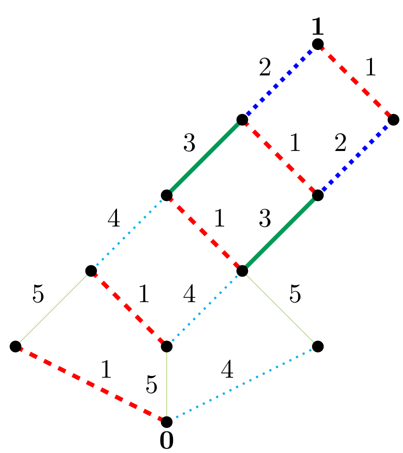

Figure 4. The S-coloring corresponding to the L-coloring of the LLD Lattice P in Figure 3. Color 2 has been subdivided into two colors (2 and 4), and color 3 has been subdivided into two colors (3 and 5). Each downward path from $\mathbf{1}$ to $\mathbf{0}$ contains one edge of each color.

We provide the characterization of LLD lattices by relating S-colorings to a specific type of L-coloring. Define a **saturated L-coloring** as an L-coloring that does not have any proper refinement that is also an L-coloring. We show that every S-coloring is an L-coloring, and every L-coloring has a refinement that is an S-coloring (which is unique up to permutation of the colors). The S in S-coloring stands for "saturated."

3.1. **Locally Confluent Processes.** L- and S-colorings are used in the study of locally confluent processes. Define a **process** $R = (P, E)$ as follows. Let P be a poset and E be the set of downward edges in the Hasse diagram of P. We say that P is the poset **generated** by process R, and R is the **underlying process** of poset P. Each element of P is a **configuration**, and each edge in E is a **move**. If u and v are configurations in P such that $u \gtrdot v$, then we say that "there is a move from u to v" or "it is possible to get from u to v in one move." If P admits an S-coloring then each color in the S-coloring of E defines a "move type."

Note from this construction that if a process R generates a poset P, then there cannot be an instance of the process that reaches the same configuration multiple times. We say that a process is **acyclic** if it satisfies this property.

A process P is **locally confluent** if for every $u, v, w \in P$ with $(u, v), (u, w) \in E$, there exists a $z \in P$ such that $(v, z), (w, z) \in E$. Local confluence is often studied in part because it implies global confluence: any finite process with a unique initial configuration (corresponding to a maximum element in P), that satisfies local confluence, must have a unique final configuration (minimum element in P). Furthermore, all complete downward paths through P must have the same length.

Furthermore, local confluence is important to us because it is closely related to the diamond property of L- and S-colorings. In fact, the diamond property on a Hasse diagram is precisely the same as local confluence on the underlying process, with additional rules about how the edges must be colored when a diamond is formed.

There do exist locally confluent processes that do not generate LLD lattices (see Figure 5). However, many of the processes that concern us are locally confluent processes that generate LLD lattices. It is often possible to prove that a locally confluent process generates an LLD lattice by showing that certain types of moves correspond to colors in an L- or S-coloring.

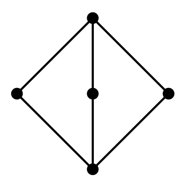

FIGURE 5. An example of a poset representing a locally confluent process, but that is not L- or S-colorable. Any coloring that satisfies the diamond property must color every edge the same color (see Figure 15). However, any coloring with just one color violates condition (1) of both L- and S-coloring.

4. S-COLORINGS

We begin by showing that a finite poset is LLD if and only if it has a maximum element and is S-colorable. We do this by relating S-colorings to L-colorings. By Theorem 3.3, a poset is an LLD lattice if and only if it is L-colorable and has a global maximum.

Lemma 4.1. *Every S-coloring is an L-coloring.*

Proof. Condition (2) is identical for both colorings. Thus, we must show that an S-coloring satisfies condition (1) of L-coloring. Suppose that a poset P has a coloring c that satisfies condition (2) of L-coloring, but not condition (1). Thus, there exist $u, v, w \in P$ and edges $(u, v), (u, w) \in E$ such that $c\big((u, v)\big) = c\big((u, w)\big)$. By condition (2), there exists a $z \in P$ such that $(v, z), (w, z) \in E$, $c\big((v, z)\big) = c\big((u, w)\big)$, and $c\big((w, z)\big) = c\big((u, v)\big)$. This implies that $c\big((u, v)\big) = c\big((u, w)\big) = c\big((v, z)\big)$, so there is a downward path from u to z containing two edges of the same color. As a result, c violates condition (1) of S-coloring. $\square$

Next we show the other direction, that every L-coloring can be refined into an S-coloring. We begin by showing that each L-coloring can be refined into a single saturated L-coloring, which is unique up to permutation of the colors.

Lemma 4.2. *Every L-coloring can be refined into a saturated L-coloring. Furthermore, if c_1 and c_2 are saturated L-colorings, then each color of c_1 is equal to a color of c_2.*

Proof. We show this by providing a construction of the unique saturated L-coloring for an arbitrary L-poset P. Given two edges (u, v) and (w, z) of the Hasse diagram of P, we say that $(u, v) \sim (w, z)$ if $(u, w) \in E$ and $(v, z) \in E$, or if $(w, u) \in E$ and $(z, v) \in E$. In other words, two edges are related by $\sim$ if they appear opposite each other in a diamond in the Hasse diagram. Let $\approx$ be the reflexive, transitive closure of $\sim$ (thus making it an equivalence relation), and define a coloring c_1 in which one color is assigned to each equivalence class of $\approx$. By definition, c_1 must meet condition (2) of L-coloring. Furthermore, if $c_1((u, v)) = i$, then under any valid L-coloring, we must have $c_1((u', v')) = i$ for all $(u', v') \in E$ such that $(u', v') \approx (u, v)$.

Thus, for any L-coloring c of P and any color c^i of c, color c^i must be a union of equivalence classes of $\approx$, or equivalently, a union of colors of c_1. If c satisfies condition (1) of L-coloring, then c_1 must also satisfy condition (1) of L-coloring, and is thus an L-coloring. As a result, c_1 is a refinement of every L-coloring, so every L-coloring can be refined into a saturated L-coloring.

Suppose that there exists some other saturated L-coloring $c_2 \neq c_1$. We have shown that c_1 must be a refinement of c_2. However, if any color of c_2 is a union of 2 or more colors of c_1, then c_2 cannot be a saturated L-coloring. Thus, each color of c_2 must also be a color of c_1, so any two saturated L-colorings of P must include the same colors. $\square$

Example 4.3. The L-coloring from Figure 3 is refined into the saturated L-coloring in Figure 4. It can be confirmed that any saturated L-coloring of P must include the same 5 colors. Note that every complete downward path through the Hasse diagram contains exactly 1 edge of each color.

Next, we provide a few lemmas to build up to the proof that any saturated L-coloring is an S-coloring. We show a correspondence between the colors in a saturated L-coloring and the join-irreducibles of the poset P. In particular, each color in a saturated L-coloring corresponds to the set of edges in E with a particular join change.

Lemma 4.4. *Let P be a finite LLD lattice with saturated L-coloring c. Then for each color c^i of the coloring c, there exists an edge $(u, v) \in E$ such that u is a join-irreducible of P, and $c\big((u, v)\big) = i$.*

Proof. Consider a color c^i of coloring c. Consider the set $U = \{u : (u, v) \in c^i\}$, the set of upper endpoints of edges in c^i. Since P is a finite poset, U, when viewed as a subposet of P, must have some minimal element u, which is the upper endpoint of edge $(u, v) \in c^i$.

Suppose for contradiction that u is not a join-irreducible. Then there must exist some $w \neq v$ such that $(u, w) \in E$. By property (2) of L-coloring, there must exist a z such that $(v, z) \in E$, $(w, z) \in E$, and $c\big((w, z)\big) = c\big((u, v)\big) = i$. Thus, we must have $w \in U$. However, $w < u$, which contradicts that u is a minimal element of U. Thus, u must be a join-irreducible of P. As a join-irreducible u can be found for any color c^i, every color must contain an edge with an upper endpoint that is a join-irreducible, as desired. $\square$

Lemma 4.5. *Let P be a finite LLD lattice with saturated L-coloring c. If $c\big((u, v)\big) = c\big((w, z)\big)$ for two edges $(u, v), (w, z) \in E$, then $J_u \backslash J_v = J_w \backslash J_z$.*

Proof. We will show that the result holds if $(u, v) \sim (w, z)$, and the result follows by transitivity. Consider edges (u, v) and (w, z) such that $(u, w), (v, z) \in E$, and $c\big((u, v)\big) = c\big((w, z)\big)$. Suppose that $J_u \backslash J_v = \{j_1\}$ and $J_v \backslash J_z = \{j_2\}$. Thus, $J_u \backslash J_z = \{j_1, j_2\}$. It is not possible to have $J_u \backslash J_w = \{j_1\}$ and $J_w \backslash J_z = \{j_2\}$, since this would imply $J_v = J_w$, which would violate the uniqueness of join-irreducible representations in LLD lattices. Thus, we have $J_u \backslash J_w = \{j_2\}$ and $J_w \backslash J_z = \{j_1\}$, which implies that $J_u \backslash J_v = J_w \backslash J_z$, as desired. Since the result holds when $(u, v) \sim (w, z)$, and since $c\big((u, v)\big)$ is an equivalence class of the transitive closure of $\sim$, the result must hold for any two edges of the same color. $\square$

Theorem 4.6. *Consider a finite LLD lattice P with saturated L-coloring c with n colors. Then for all $i \in [n]$, there exists a join-irreducible j of P such that $c^i = \big\{(u, v) : J_u \backslash J_v = \{j\}\big\}$.*

Proof. Given a join-irreducible j, let $E_j = \big\{(u, v) \in E : J_u \backslash J_v = \{j\}\big\}$.

Consider a color c^i from saturated L-coloring c. By Lemma 4.4, c^i contains an edge (j, k) such that j is a join-irreducible of P. The join change of edge (j, k) is j, so by Lemma 4.5, all edges of c^i have join change j, so $c^i \subseteq E_j$.

Similarly, for any color $c^{i'} \neq c^i$, $c^{i'}$ must contain an edge (j', k') such that j' is a join-irreducible, and $j' \neq j$. All edges of $c^{i'}$ must have join change $j' \neq j$, so $E_j \subseteq c^i$.

Thus, for any color c^i from a saturated L-coloring c, we must have $c^i = E_j$ for some join-irreducible j, as desired. $\square$

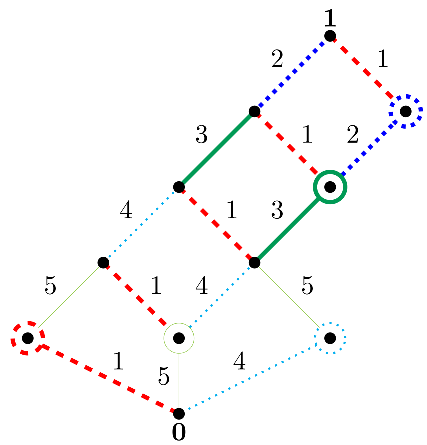

FIGURE 6. In the LLD lattice from Figure 3, each join-irreducible is circled using the color of its lone downward edge. Note that there is exactly one join-irreducible of each color. For each color, all edges of that color have a join change equal to the corresponding join-irreducible.

In addition to being useful in its own right, Theorem 4.6 allows us to complete our correspondence between L-colorings and S-colorings.

Lemma 4.7. *Every saturated L-coloring is an S-coloring.*

Proof. As condition (2) of L-colorings and S-colorings are the same, we must prove that every saturated L-coloring satisfies condition (1) of S-coloring.

Consider a finite LLD lattice P with saturated L-coloring c, and consider a complete downward path $S = \big((v_0, v_1), (v_1, v_2), \dots (v_{n-1}, v_n)\big)$, where $v_0 = \mathbf{1}$ and $v_n = \mathbf{0}$. The join changes of the n edges in this path must be distinct, so by Theorem 4.6, the n edges must have n distinct colors under coloring c. This must be true for every complete downward path, so condition (1) of an S-coloring is satisfied, as desired. $\square$

Lemma 4.8. *A finite poset is an LLD lattice if and only if it has a maximum element and admits an S-coloring.*

Proof. By lemmas 4.1 and 4.7, a finite poset with a maximum element is S-colorable if and only if it is L-colorable. Since a finite poset with a maximum element is L-colorable if and only if it is LLD, a finite poset with a maximum element is S-colorable if and only if it is LLD. $\square$

Lemma 4.9. *Every S-coloring is a saturated L-coloring.*

Proof. Consider a finite LLD lattice P with S-coloring c and (unique) saturated L-coloring c'. By Lemma 4.1, c is an L-coloring. Thus, by Lemma 4.2, each color in c is a union of colors in c'.

Suppose for contradiction that c is not a saturated L-coloring. Thus, there must be some color c^i of c that is a union of two or more colors in c'. Suppose that c'^{i_1} and c'^{i_2} are two colors of c' such that $c'^{i_1} \cup c'^{i_2} \subseteq c^i$. By 4.6, there exist join-irreducibles j_1 and j_2 such that c'^{i_1} consists of all edges with join change j_1, and c'^{i_2} consists of all edges with join change j_2.

Now, consider a complete downward path $S = \big((v_0, v_1), (v_1, v_2), \dots (v_{n-1}, v_n)\big)$, where $v_0 = \mathbf{1}$ and $v_n = \mathbf{0}$. As $\mathbf{1}$ is the join of all join-irreducibles of P, and $\mathbf{0}$ is the empty join of join-irreducibles, there must exist some edge $e \in S$ with join change j_1, and some edge in $e' \in S$ with join change j_2. However, e and e' are both of color c^i in the S-coloring c. This violates condition (1) of S-coloring, so every S-coloring is a saturated L-coloring, as desired. $\square$

Lemmas 4.7 and 4.9 combine to give the following result.

Theorem 4.10. *Consider a finite LLD lattice P and a coloring c. Then c is an S-coloring if and only if it is a saturated L-coloring.*

We thus have a correspondence between saturated L-colorings and S-colorings, in which each color corresponds to the join change of all of the edges of that color. This is summarized in the following result.

Theorem 4.11. *Each LLD lattice P has a unique S-coloring up to isomorphism. In this S-coloring, each edge is colored according to its join change.*

Proof. By Theorem 4.10, a coloring is an S-coloring if and only if it is a saturated L-coloring, and by Theorem 4.6, P has a unique saturated L-coloring up to isomorphism. In this coloring, each edge is colored according to its join change. □

4.1. **The Color Poset.** A consequence of Theorem 4.6 is that for any S-coloring c and any complete downward path S in a finite LLD lattice, S must contain one edge of each color in c. Furthermore, for any elements $u > v$, any two downward paths from u to v must contain edges of the same colors, in a possibly different order.

Thus, it is well-defined to consider whether one color "occurs before" another color on a path from the top to the bottom. As a result, the color poset is well-defined: given a finite LLD lattice P and S-coloring c with colors c^1 and c^2, we say that $c^1 \geq c^2$ in the **color poset** if there is no complete downward path in which the edge of color c^2 occurs before the edge of color c^1. Our main result of this subsection is that the colors of the S-coloring of P induce the same partial ordering as join-irreducibles corresponding to each color.

Theorem 4.12. *The color poset of an LLD lattice is isomorphic to its poset of join irreducibles.*

Proof. Consider a finite LLD lattice P, with join-irreducibles u and w, which are upper endpoints of the edges (u, v) and (w, z) of the Hasse diagram of P. We want to show that $c\big((u, v)\big) \geq c\big((w, z)\big)$ in the color poset, if and only if $u \geq w$.

First suppose $c\big((u, v)\big) \geq c\big((w, z)\big)$. No downward path from $\mathbf{1}$ to u may contain an edge of color $c\big((w, z)\big)$. Thus, every downward path from u to $\mathbf{0}$ contains an edge of color $c\big((w, z)\big)$. Thus, starting at u, we can follow downward edges of colors other than $c\big((w, z)\big)$ until no edges of colors other than $c\big((w, z)\big)$ are available. The resulting element of the poset must be w, the join-irreducible from which the only downward edge is of color $c\big((w, z)\big)$. Since there is a downward path from u to w, we have $u \geq w$ as desired.

Instead, suppose that $c\big((u, v)\big) \not\geq c\big((w, z)\big)$. Consider some complete downward path $s = (e_1, e_2, \ldots e_n)$ in which the edge of color $c\big((w, z)\big)$ occurs before the edge of color $c\big((u, v)\big)$. Thus, $c(e_j) = c\big((w, z)\big)$ and $c(e_k) = c\big((u, v)\big)$ for some $k > j$. Create a new complete downward path s' as follows: Let the first $k - 1$ edges in s' be the same as the first $k - 1$ edges in s. From there, instead of following edge e_k, follow edges of colors other than $c\big((u, v)\big)$ until no edges of colors other than $c\big((u, v)\big)$ are available. The resulting vertex must be u, from which the only downward edge is (u, v). From there, follow edge (u, v) and then follow any downward path from v to $\mathbf{0}$.

Thus, there is a downward path from $\mathbf{1}$ to u containing an edge of color $c\big((w, z)\big)$. Since an element of the poset uniquely corresponds to the colors of downward edges needed to reach it from $\mathbf{1}$, every downward path from $\mathbf{1}$ to u must contain an edge of color $c\big((w, z)\big)$, and no downward path from u to $\mathbf{0}$ may contain an edge of color $c\big((w, z)\big)$. In particular, no downward path from u to $\mathbf{0}$ may contain edge (w, z). Since (w, z) is the only downward edge from w, there can be no downward path from u to w, so $u \not\geq w$, as desired. □

Example 4.13. In the poset P from Figure 6, the dark blue join-irreducible is greater than the light green join-irreducible. This is consistent with the fact that every complete downward path must go through a dark blue edge before it goes through a light green edge, which implies that the color dark blue (number 2) is greater than the color light green (number 5) in the color poset.

We will apply Theorem 4.12 to a variety of processes and other types of lattices, where the colors represent certain types of moves. The following lemmas help to analyze what the colors represent in these processes. First, we show that, given an upper and lower endpoint in an LLD lattice, it is possible to transform one downward path into another by "switching pairs of consecutive edge colors" one at a time.

Given two downward paths $s = (e^1, e^2, \ldots e^n)$ and $s' = (e'^1, e'^2, \ldots e'^n)$, we say that s and s' are related by an **edge pair swap** if there exists an index i such that the following holds:

- $e^h = e'^h$ for all $h < i$
- $c(e^i) = c(e'^{i+1})$
- $c(e^{i+1}) = c(e'^i)$
- $e^h = e'^h$ for all $h > i+1$

Using this notation, we obtain the following result:

Lemma 4.14. *Consider a finite LLD lattice P with S-coloring c. Consider vertices $u, v \in P$ where $u \geq v$, and two downward paths $s_0 = (e_0^1, e_0^2, \ldots e_0^n)$ and $s_f = (e_f^1, e_f^2, \ldots e_f^n)$ from u to v. Then there exists a sequence of downward paths from u to v $(s_0, s_1, s_2, \ldots s_k = s_f)$ $\big($where for each j, sequence $s_j = (e_j^1, e_j^2, \ldots e_j^n)\big)$, such that for each $1 \leq j \leq k$, s_{j-1} and s_j are related by an edge pair swap.*

Proof. Induct on n. If v can be reached from u in one edge, then there is only one downward path from u to v, so the result holds for $n = 1$.

Now, assume that the result holds for all downward paths of length $n-1$. Consider edge e_f^1, the first edge of s_f, and consider the edge e_0^k of s_0 such that $c(e_0^k) = c(e_f^1)$. If $k = 1$, then s_0 and s_f both contain a path from the lower endpoint of e_f^1 to v of length $n-1$. Thus, by our inductive assumption, s_0 can be transformed into s_f through a series of edge pair swaps.

Otherwise, assume that $k \neq 1$. There is a downward edge of color $c(e_f^1)$ starting at u. Thus, by condition 2 of S-coloring, any downward path starting from u that does not contain an edge of color $c(e_f^1)$, must lead to a vertex that is the start point of an edge of color $c(e_f^1)$. In particular, an edge of color $c(e_f^1)$ must be available after following the downward path $(e_0^1, e_0^2, \ldots e_0^{k-2})$ from u. Thus, edges of colors $c(e_0^k)$ and $c(e_0^{k-1})$ must both be available from the end of this path, so by the L-coloring property, the element of P reachable by following the edge of color $c(e_0^{k-1})$ and then the edge of color $c(e_0^k)$, is also reachable by following the edge of color $c(e_0^k)$ and then the edge of color $c(e_0^{k-1})$. By performing an edge color swap that reverses the order of colors $c(e_0^k)$ and $c(e_0^{k-1})$, we create a new downward path s_1 in which the edge of color e_f^1 appears earlier in the sequence than it does in path s_0.

Performing $k-1$ edge pair swaps involving edges of color e_f^1 results in a sequence s_{k-1} such that $c(e_{k-1}^1) = c(e_f^1)$. Again from our inductive assumption, s_{k-1} can be transformed into s_f through a series of edge pair swaps, so s_0 can be transformed into s_f through edge pair swaps. $\square$

Lemma 4.15. *Given a finite LLD lattice P with S-coloring c, colors c^1 and c^2 are incomparable in the color poset of c if and only if there exist vertices $u, v, w \in P$ such that:*

- $(u, v), (u, w) \in E$.
- $c\big((u, v)\big) = c^1$
- $c\big((u, w)\big) = c^2$.

Proof. If such vertices u, v, and w exist, then by condition (2) of S-coloring, there exists a vertex z and edges (v, z) and (w, z) such that $c\big((v, z)\big) = c^2$ and $c\big((w, z)\big) = c^1$. As a result, there are downward paths from u to z (and thus from $\mathbf{1}$ to $\mathbf{0}$) where colors c^1 and c^2 occur in either order. Thus, c^1 and c^2 must be unrelated.

Now, suppose that c^1 and c^2 are unrelated, so that there exist two different complete downward paths s_1 and s_2, where s_1 contains its edge of color c^1 before its edge of color c^2, but s_2 contains its edge of color c^2 before its edge of color c^1. By Lemma 4.14, it is possible to transform s_1 into s_2 through successive operations of replacing two edges with two edges with the same edge colors in the opposite order. At some point, we must reach intermediate states s_1' and s_2', where the edge of color c^1 appears immediately before the edge of color c^2 in s_1', the edge of color c^2 appears immediately before the edge of color c^1 in s_2', and all other edges are the same in s_1' and s_2'. Thus, there exists some k such that $c((s_1')_k) = 1$, $c((s_1')_{k+1}) = 2$, $c((s_2')_k) = 2$, and $c((s_2')_{k+1}) = 1$. Thus, starting from $\mathbf{1}$ and following the first $k-1$ edges of s_1' results in a vertex from which downward edges of colors c^1 and c^2 are both present, as desired. $\square$

The following lemma allows us to evaluate individual L-colorings in terms of the corresponding S-coloring.

Lemma 4.16. *Consider a finite LLD lattice P with an L-coloring c_1 and S-coloring c. Consider a single color c_1^i of c_1. All of the colors of c that are subsets of c_1^i form a chain in the color poset.*

Proof. Since c_1 is an L-coloring, there are no two edges of color c_1^i with the same top vertex. By Lemma 4.15, any two colors of c that are subsets of c_1^i must be comparable in the color poset, so they form a chain. $\square$

Lemma 4.16 leads directly to the following corollary, which allows each color of an S-coloring to be interpreted as "the j^{th} move of a given type."

Corollary 4.17. *Consider a finite LLD lattice P with an L-coloring c_1 and S-coloring c. Consider a single color c^{i_1} of c, which is a subset of color $c_1^{i_2}$ of c_1. Then there exist nonnegative integers j_1 and j_2 such that for every complete downward path through the Hasse diagram of P, the edge of color c^{i_1} occurs after j_1 other edges of color $c_1^{i_2}$, and before j_2 other edges of color $c_1^{i_2}$*

Example 4.18. In poset P from Figure 4, the dark blue color is greater than the light blue color, and the dark green color is greater than the light green color in the color poset. The two blue colors (and, respectively, the two green colors) are guaranteed to be comparable in the color poset by Lemma 4.16, because both blue colors were part of the original, single blue color from Figure 3, while both green colors were part of the original green color.

Lemma 4.16 and Corollary 4.17 are relevant to many applications of LLD lattices, in which LLD lattices are used to represent processes. There is often a natural L-coloring in which each color corresponds to a certain type of move, and Corollary 4.17 tells us that each color in the S-coloring corresponds to the "j^{th} move of a given type."

4.2. **Shortest Paths and S-Colorings.** A consequence of Theorem 4.12 is that for any elements $u \geq v$ of an L-colored lattice P, any sequence of downward edges from u to v must consist of the same collection of distinct colors in a possibly different order. We can now generalize this downward path result to a result on *any* sequence of (upward or downward) edges between any two elements u and v. This idea has applications to many of the processes we care about, most notably domino tilings and transformations between them.

Consider a finite LLD lattice P with maximum element $\mathbf{1}$, minimum element $\mathbf{0}$, and S-coloring c with n distinct colors $c^1, \dots c^n$. Given an element u, define the **color vector** of u, $V(u) = \big(V_1(u), V_2(u), \dots V_n(u)\big)$ as follows. For each i from 1 to n, let $V_i(u)$ be 1 if an edge of color c^i is included in any downward path from $\mathbf{1}$ to u, and 0 otherwise. Note that as a consequence of Theorem 4.12, this color vector is uniquely defined for each element of P.

Given two elements u and v and a path $s = (e_1, e_2, \dots e_k)$ from u to v in the Hasse diagram of P, define the **color vector** $V(s) = \big(V_1(s), V_2(s), \dots V_n(s)\big)$ as follows. For each i from 1 to n, $V_i(s)$ is equal to the number of downward edges of color c^i minus the number of upward edges of color c^i in s. We show that this color vector is the same for all paths from u to v.

Lemma 4.19. *Consider a finite LLD lattice P with S-coloring c, elements u and v, and a path s from u to v. If $V(s)$ is the color vector of s defined using coloring c, then $V(s) = V(v) - V(u)$.*

Proof. We show that $V(v) = V(u) + V(s)$ by inducting on the length of s. If there are no edges in s, then $v = u$, so $V(v) = V(u)$. Now, suppose that the equality holds for all sequences of length n. Consider a sequence $s = (s_1, s_2, \dots s_{n+1})$ of length $n+1$ from u to v, with subsequence $s' = (s_1, s_2, \dots s_n)$ of length n from u to v'. By our inductive assumption, we have $V(v') = V(u) + V(s')$. Then consider edge s_{n+1}. If s_{n+1} is a downward edge of color c^i, then $V(s) = V(s') + e_i$, where e_i is the unit vector consisting of all 0's, except for a 1 in the i^{th} coordinate. Furthermore, configuration v is reachable from $\mathbf{1}$ by first reaching v', and then following an additional downward edge of color c^i, so $V(v) = V(v') + e_i$. As a result, we get $V(v) = V(v') + e_i = V(u) + V(s') + e_i = V(u) + V(s)$ as desired.

If s_{n+1} is an upward edge of color c^i, then $V(s) = V(s') - e_i$. Furthermore, v' is reachable from $\mathbf{1}$ by first reaching v, and then following an additional downward edge of color c^i, so $V(v') = V(v) + e_i$, or equivalently, $V(v) = V(v') - e_i$. As a result, we get $V(v) = V(v') - e_i = V(u) + V(s') - e_i = V(u) + V(s)$ as desired. □

The fact that the color vector is the same for every sequence s connecting u to v allows us to define the **color vector** of a pair of elements $V(u, v)$ as $V(s)$ for any path s from u to v.

Lemma 4.20. *Consider a finite LLD lattice P with S-coloring c, and elements u and v. If $V(u, v)$ is the color vector of u, v defined using coloring c, then all entries of $V(u, v)$ are in $\{-1, 0, 1\}$.*

Proof. By Theorem 4.12, any downward path from $\mathbf{1}$ to any element u contains at most one edge of each color. As a result, all entries of $V(u)$ and $V(v)$ are in $\{0, 1\}$. Subtracting two vectors with entries in $\{0, 1\}$ yields a vector with entries in $\{-1, 0, 1\}$. □

Theorem 4.21. *Consider a finite LLD lattice P with S-coloring c, and elements u and v. If $V(u,v)$ is the color vector of u,v defined using coloring c, then the shortest path from u to v consists of $||V(u,v)||_{L^1}$ edges. Furthermore, there exists a path of length $||V(u,v)||_{L^1}$ in which all downward edges precede all upward edges.*

Proof. Consider the vector V_{max}, the pointwise maximum of $V(u)$ and $V(v)$. We will show that there exists an element $w \in P$ such that $u \geq w$, $v \geq w$, and such that $V(w) = V_{max}$. Thus, we can achieve the desired bound by moving from u to w through a sequence of downward edges, and then to v through a sequence of upward edges.

We show that there exists an element w such that $u \geq w$ and $V(w) = V_{max}$. We prove that from u, there exists a downward edge of some color c^i such that $V(v)_i = 1$ and $V(u)_i = 0$. Following this edge, and then repeatedly applying the argument to the resulting element will eventually yield w. Suppose that there exists a path from $\mathbf{1}$ to u consisting of edges $e_1, e_2, \dots e_k$ with corresponding colors $c^1, c^2, \dots c^k$, and there exists a path from $\mathbf{1}$ to v consisting of edges $e'_1, e'_2, \dots e'_l$ with corresponding colors $(c^1)', (c^2)', \dots (c^l)'$. Define s_1 to be the sequence $(c^1, c^2, \dots c^k)$ and s_2 to be the sequence $\big((c^1)', (c^2)', \dots (c^l)'\big)$. Let $(c^n)'$ be the first color in s_2 that does not appear in s_1. Thus, $(c^1)', (c^2)', \dots (c^{n-1})'$ all appear in s_1. By condition (2) of S-coloring, s_1 can be rearranged into a sequence s'_1 that still starts at $\mathbf{1}$, ends at u, and contains the same edge colors as s_1, but which begins with edge colors $(c^1)', (c^2)', \dots (c^{n-1})'$. A downward path starting at $\mathbf{1}$ and following edges of colors $(c^1)', (c^2)', \dots (c^{n-1})'$ in order gives some configuration u_n such that:

- There exists a downward path from u_n to u.
- There is an edge of color $(c^n)'$ with upper endpoint u_n

Since $(c^n)'$ does not appear in s_1, by the L-coloring property, there is also an edge of color $(c^n)'$ with upper endpoint u.

Starting at u and following the downward edge of color $(c^n)'$ reaches a configuration u', reachable from u through one downward edge, such that $V(u')$ is equal to $V(u)$, except for a single 0 replaced with a 1 in the place corresponding to color $(c^n)'$. As $(c^n)'$ is in s_2 but not s_1, the corresponding entry in $V(v)$ (and thus V_{max}), is equal to 1, while the corresponding entry in $V(u)$ is equal to 0. Thus, $||V(u) - V_{max}||_{L^1} - ||V(u') - V_{max}||_{L^1} = 1$. Since the initial distance between $V(u)$ and V_{max} must be finite, we can show by induction that by repeatedly following downward edges in this manner, an element w whose color vector is V_{max} is eventually reached. We can show that the same is possible from v, and since a configuration is uniquely determined by its color vector, this same configuration w must be a lower bound for u and v.

Thus, it is possible to go from u to w through a sequence of downward edges, and then to v through a sequence of upward edges with the total number of moves equal to $||V(u,v)||_{L^1}$.

Since following one edge changes a single entry in the color vector by exactly 1, $||V(u,v)||_{L^1}$ is the smallest number of edges possible to get from u to v. $\square$

Example 4.22. In Figure 7, vertices u and v are labeled in the LLD poset from Figure 4. Every shortest path from u to v contains one red edge, one light blue edge, and one dark green edge.

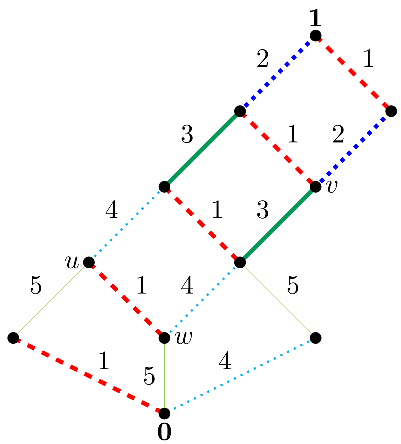

FIGURE 7. Every shortest path from u to v contains one red edge, one light blue edge, and one dark green edge. In every shortest path, the red edge is traversed downward, while the light blue and dark green edges are traversed upward. There exists a path through vertex w in which the downward red edge is traversed before the two upward edges.

This framework allows us to deal with many types of processes that generate LLD lattices, in which downward edges represent moves between configurations in a given state space, and colors represent certain classes of moves. We can prove that a process forms an LLD lattice by showing that it meets the L-coloring property, usually by associating each color with a certain class of moves (e.g. a color corresponding to all firing moves at a given site in chip-firing). We can then refine this L-coloring into a saturated L-coloring, or equivalently, an S-coloring. If it is never possible to have two moves of the same class available from the same configuration (e.g. there can be at most one firing move available at any given site from a given configuration in chip-firing), then there is a further interpretation of the colors in the S-coloring. Each corresponds to "the j^{th} move of a given class" (the first firing move at site i, the second firing move at site i, etc.) by Corollary 4.17.

In some settings, including domino tilings, there is little distinction between which moves are considered "forward" (corresponding to downward edges in P), and which moves are considered "backward" (corresponding to upward edges). In contexts without that distinction, our results on shortest paths are still useful, as those results do not require us to only consider paths going in a single direction.

5. Chip-Firing

Now that we have the framework and main results on using colorings in LLD lattices, we can apply the results of Section 3 to processes that generate LLD lattices. In particular, we use our methods to provide a new proof of a key sorting result in labeled chip-firing.

Consider an acyclic process R generating a poset P. Each vertex in P corresponds to a configuration of the process, and each edge in the Hasse diagram of P corresponds to a possible move between two configurations in the process. A color in the color poset of P corresponds to some type of move, which often has a natural interpretation in the underlying process.

We first consider chip-firing. In a chip-firing process, a collection of indistinguishable chips are placed at the nodes of a graph. If a node has at least as many chips as it has outgoing edges, it can "fire" by sending one chip along each outgoing edge to the node's corresponding neighbors. The process terminates if no site has enough chips to fire. Throughout this section, we assume that the chip-firing process begins with an initial configuration u_s and terminates with a final configuration u_f.

Define a **chip configuration** on a graph $G = (V, E)$ to be a $\mathbb{N}$-valued function on V, designating the number of chips at each $v \in V$. On a terminating chip-firing process with initial configuration u_0, the **configuration poset** P is defined so that configuration $u_1 \geq u_2$ in P if it is possible to get from u_1 to u_2 through some sequence of firing moves. An edge in the Hasse diagram of P corresponds to a single firing of exactly $deg(v)$ chips at a single vertex v.

For any terminating chip-firing process, there is a natural coloring for the edges: each color corresponds to all firing moves at a single vertex. Since there can be at most one available move at each site from a given configuration, this coloring must meet condition (1) for L-coloring. As chip-firing is locally confluent (any two available firing moves can be performed in either order), this coloring also meets condition (2). Thus, this coloring is an L-coloring.

Consider a vertex v. By Lemma 4.16 and Corollary 4.17, the colors in the S-coloring corresponding to firing moves at v form a chain in the color poset. We are thus able to interpret each color in the S-coloring as "the j^{th} move at site k" for some fixed j.

The color poset in chip-firing is crucial to labeled chip-firing, in which the local confluence of chip-firing is used to prove confluence for a form of chip-firing with distinguishable chips. In labeled chip-firing, n chips are placed at the origin of a 1D grid graph. Here, the individual chips are given distinct labels from 1 to n. When two chips are fired, we impose the additional rule that the chip with the smaller label is sent to the left, and the chip with the larger label is sent to the right. Hopkins, McConville and Propp [9] and Klivans and Liscio [11] prove the following result:

Theorem 5.1 ([9], Theorem 11 and [11], Theorem 2.10)**.** *The labeled chip-firing process with $2m$ chips at the origin terminates in a final configuration in which all chips are in sorted order from left to right.*

This is a notable result, as it is a case in which global confluence exists without local confluence.

The proof of Klivans and Liscio is based on an analysis of the "move poset" of unlabeled chip-firing on the line. By showing that certain moves (of the form "j^{th} move at site k") near the end of the process must take place in a certain order, and that those moves must take place with just enough chips to fire, we can show that each labeled chip must be restricted to specific places at specific times throughout the process.

We have the following result relating the move poset in chip-firing to the join-irreducibles of the configuration poset:

Theorem 5.2. *Consider a terminating chip-firing process on a graph G with initial configuration u. Then the move poset of the resulting process is isomorphic to the poset of join-irreducibles of the poset of configurations reachable from u.*

Proof. By Lemma 4.16 and Corollary 4.17, each color c^i contains all edges corresponding to a move of the form "the j^{th} move at site k" for some j and k. Thus, each color c^i corresponds precisely to a move in the move poset. Furthermore, color $c^i \geq c^{i'}$ in the color poset if and only if the move corresponding to c^i must occur before the move corresponding to $c^{i'}$. Thus, the color poset is isomorphic to the move poset. By Theorem 4.12, the color poset is isomorphic to the poset of join-irreducible configurations in the configuration poset, so the move poset is isomorphic to the poset of join-irreducibles of the configuration poset. □

By viewing the move poset as the poset of join-irreducibles, or equivalently, as the color poset on the S-coloring of the poset of chip configurations, we can simplify the proof of this sorting result, and also make the proof more adaptable to other related problems.

Klivans and Liscio prove sorting in labeled chip-firing using a 3-step process [11]:

(1) Prove that there is a diamond of firing moves at the bottom of the Hasse diagram of the move poset. Each firing move in this diamond takes place with exactly 2 chips present.
(2) Prove bounds on how far to the left the largest chips can go throughout the process, and how far to the right the smallest chips can go.
(3) Combine steps 1 and 2 to restrict the position of each chip during each firing move in the diamond, and ultimately to restrict the position of each chip at the end of the process.

As the move poset is isomorphic to the poset of join-irreducibles of the configuration poset, we can reprove step 1 by proving an ordering on join-irreducible configurations, instead of proving an ordering on firing moves. This simplifies the proof, as it is often easier to prove that one configuration *can* be reached from another than to prove that one move *must* occur before another.

We make use of the following result from [1]:

Theorem 5.3. *The chip-firing process with $2m$ chips at the origin terminates at the final configuration with a single chip at every position from $-m$ to -1, and from 1 to m.*

This gives us enough information to complete step (1) of the sorting proof. In the original proof from [11], an additional result was needed, which proved the number of firing moves that must occur at each site throughout the process. Here, that is replaced with Corollary 4.17, since Corollary 4.17 is enough to guarantee that "the j^{th} to last move at site k" is well-defined.

Define move k_j to be the j^{th}-to-last firing move at site k (which is well defined because of Corollary 4.17). By using the fact that the color poset is equivalent to the poset of join-irreducibles, we can simplify the proof of the following lemma (equivalent to Lemma 2.6 form [11]), removing two layers of induction from the proof. Given a firing move k_j, we define $J(k_j)$ to be the join-irreducible corresponding to move k_j. This is the minimum element in the poset of configurations at which a move of type k_j can occur, or equivalently, the unique configuration from which a move of type k_j is the only one available.

Lemma 5.4. *For all $0 \leq j \leq m$, move 0_j must take place before moves 1_{j-1} and -1_{j-1}. For $1 \leq k \leq m$ and $0 \leq j \leq m-k$, move k_j must take place before moves $(k+1)_{j-1}$ and $(k-1)_j$. For $-m \leq k \leq -1$ and $0 \leq j \leq m-|k|$, move k_j must take place before moves $(k-1)_{j-1}$ and $(k+1)_j$.*

Proof. By Theorem 4.12, a move $(k_1)_{j_1}$ being required to take place before $(k_2)_{j_2}$ is equivalent to $J\big((k_2)_{j_2}\big)$ being reachable from $J\big((k_1)_{j_1}\big)$. As a result, we will show that $J(1_{j-1})$ is reachable from $J(0_j)$, $J(-1_{j-1})$ is reachable from $J(0_j)$, etc.

We will show that for $0 \leq j \leq m-|k|$, $J(k_j)$ consists of 2 chips at site k, 0 chips at site $k+j$ $(k \geq 0)$ or $k-j$ $(k \leq 0)$, 0 chips at site $-j$ $(k \geq 0)$ or j $(k \leq 0)$, and 1 chip at all other sites from $-m$ to m.

Assume that $k \geq 0$. The proof for $k \leq 0$ is identical, but with the signs of the sites changed.

A configuration of the specified form is reachable through the following firing moves at each site:

- all moves at sites $k+j$ through m and $-m$ through $-j$
- all but the last j moves at sites 0 through k
- all but the last $k+j-x$ moves at sites $x=k+1$ through $x=k+j-1$
- all but the last $j+x$ moves at sites $x=-j+1$ through $x=-1$

We show that these firing moves produce the desired configuration by comparing the number of firing moves at each site and its neighbors to the total number of firing moves throughout the process.

- At sites $x > k+j$ and $x < -j$, all moves have occurred at site x and its neighbors, so the number of chips at site x in configuration $J(k_j)$ is equal to the number of chips (1 chip) at site x in the final configuration.
- At sites $x = k+j$ and $x = -j$, all moves have occurred at site x and one of its neighbors, but one move has not yet occurred at its other neighbor. As a result, site x has one fewer chip (0 chips) in configuration $J(k_j)$ than site x has in the final configuration.
- If $k \neq 0$, the number of firing moves yet to occur at the origin exceeds the average number of remaining firing moves at its neighbors by $\frac{1}{2}$, so the origin has one more chip (1 chip) in configuration $J(k_j)$ than the origin has in the final configuration. If $k = 0$, then the number of firing moves yet to occur at the origin exceeds the average number of remaining firing moves at its neighbors by 1, so the origin has 2 more chips (2 chips) in configuration $J(k_j)$ than the origin has in the final configuration.
- if $k \neq 0$, the number of firing moves yet to occur at site k exceeds the average number of remaining firing moves at its neighbors by $\frac{1}{2}$, so site k has one more chip (2) in configuration $J(k_j)$ than site k has in the final configuration.
- At all other sites x, the number of firing moves yet to occur at site x is equal to the average number of remaining firing moves at its neighbors, so site x has as many chips (1) in configuration $J(k_j)$ as site x has in the final configuration.

As a result, each join-irreducible takes the form specified above. Now consider moves k_j and $(k+1)_{j-1}$. $J(k_j)$ consists of 2 chips at site k, 0 chips at site $k+j$, 0 chips at site $-j$, and 1 chip at all other sites from $-m$ to m. $J\big((k+1)_{j-1}\big)$ consists of 2 chips at site $k+1$, 0 chips at site $k+j$, 0 chips at site $-j+1$, and 1 chip at all other sites from $-m$ to m. It is possible to get from $J(k_j)$ to $J\big((k+1)_{j-1}\big)$ by firing successively at every site from k down to $-j+1$, so $J\big((k+1)_{j-1}\big)$ is reachable from $J(k_j)$. Similar arguments apply for showing comparable cases for each join-irreducible and its descendants.

Because the corresponding join-irreducibles satisfy the desired reachability constraints, the moves of the poset satisfy the corresponding order, as desired. □

The proof of Lemma 5.4 gives us the following corollary:

Corollary 5.5. *Every move k_j referenced in Lemma 5.4 must take place with exactly 2 chips present at site k.*

This completes step (1) of the proof of sorting in labeled chip-firing. From here, steps (2) and (3) can be performed as they were in [11] to complete the proof of sorting.

Remark 5.6. While color poset-based methods are used here to reprove a result on labeled chip-firing, these methods are used to prove new sorting results as well. In [12], these methods are used to classify every weakly sorted configuration on the line as either "sorting" or "non-sorting." No other proof of the classification result is currently known. Sorting for other initial weights in labeled chip-firing of type A_{n-1} (discussed in Section 6 of this paper) is a corollary of the classification theorem.

6. Root System Central-Firing

In addition to simplifying the proof of sorting for labeled chip-firing, the results here can be used to prove new confluence results. Galashin, Hopkins, McConville, and Postnikov view labeled chip-firing as a "Type A" case of a more general type of root system central-firing. We present some preliminaries on root system central-firing, and we state conjectures from [8] on which initial weights sort for types central-firing of types A, B, C, and D.

For more information on root systems, see [10], [5], or [3]. For more information on central-firing, see [8] or [12].

6.1. **Central-Firing.** We define a process on the weights of a root system Φ. Given a root system Φ with weight lattice P, and two weights $\lambda, \lambda' \in P$, we say that λ' is **strongly reachable** from λ if $\lambda' = \lambda + \alpha$ for some positive root α such that $\langle \lambda, \alpha^\vee \rangle = 0$. λ' is **reachable** from λ if there exists some sequence of weights $(\lambda = \lambda^0, \lambda^1, \dots \lambda^k = \lambda')$ such that for each $1 \le i \le k$, λ^i is strongly reachable from λ^{i-1}. Our choices of simple roots, and the resulting positive roots, for each Coxeter type, are shown in Table 1.

Type	A_{n-1}	B_n	C_n	D_n
simple roots $i \in [n]$	$\alpha_i = e_{i+1} - e_i$	$\alpha_1 = e_1$ $\alpha_i = e_i - e_{i-1}$, $i \ge 2$	$\alpha_1 = 2e_1$ $\alpha_i = e_i - e_{i-1}$, $i \ge 2$	$\alpha_1 = e_1 + e_2$ $\alpha_i = e_i - e_{i-1}$, $i \ge 2$
positive roots $i < j \in [n]$	$e_j - e_i$	$e_j \pm e_i$ e_i	$e_j \pm e_i$ $2e_i$	$e_j \pm e_i$
chip positions in ω_n	N/A	all chips at $\frac{1}{2}$	all at 1	all at $\frac{1}{2}$
chip positions in ω_{n-1}	chip 1 at origin chips 2 to n at 1	1 at origin 2 to n at 1	1 at origin 2 to n at 1	1 at $-\frac{1}{2}$ 2 to n at $\frac{1}{2}$
chip positions in other ω_i	1 to $n-i$ at origin $n-i+1$ to n at 1	1 to $n-i$ at origin $2n-i+1$ to n at 1	1 to $n-i$ at origin $n-i+1$ to n at 1	1 to $n-i$ at origin $n-i+1$ to n at 1
operations used	(a)	(a), (b), (d)	(a), (c), (d)	(a), (d)

Table 1. A summary of the simple and positive roots, fundamental weights, and types of chip-firing operations allowed in root system central-firing of types A_{n-1}, B_n, C_n, and D_n.

Given an initial weight ω, define the configuration poset P_ω to be the set of weights reachable from ω, ordered by reachability (so that $\lambda \ge \lambda'$ if λ' is reachable from λ). The initial weight ω is usually chosen to be a fundamental weight. Define the **central-firing process** on Φ with initial weight ω to be $R(\Phi, \omega) = \Big(P_\omega, E(P_\omega)\Big)$, where $E(P_\omega)$ is the cover relation of P_ω. Galashin et al. show that P_ω is finite for any Φ and ω [8].

A weight $\lambda = (\lambda_1, \lambda_2, \dots \lambda_n)$ is often thought of as representing a configuration of n chips labeled from 1 to n, where chip 1 is at position λ_1, chip 2 is at position λ_2, $\dots$ and chip n is at position λ_n. The chip configurations corresponding to fundamental weights $\omega_1, \omega_2, \dots \omega_n$ of each classical root system are shown in Table 1.

The **moves** of a central-firing process are pairs of weights (λ, λ') such that λ covers λ' in P_ω. These moves correspond to four types of operations on the chips:

(a) For $i < j$, if chips i and j are in the same position, move chip i one step to the left, and chip j one step to the right.
(b) For $i \in [n]$, if chip i is at the origin, move it one step to the right.
(c) For $i \in [n]$, if chip i is at the origin, move it two steps to the right.
(d) For $i \neq j$, if chips i and j are in opposite positions ($v_i = -v_j$), move both chips one step to the right.

Galashin et al. [8] show that for any weights λ and λ' such that $\lambda \geq \lambda'$, it is possible to get from λ to λ' by applying some combination of (a), (b), (c), and (d) operations, depending on the Coxeter type of the root system being used. The allowed operations for each classical root system are shown in Table 1. Because every move can be expressed as one of these four types of operations, we will use the term "move" instead of "operation" throughout this section.

Central-firing of type A_{n-1} is equivalent to labeled chip-firing with n chips. Note that our choices of simple and positive roots are different from the choices made in [8]. We choose our simple and positive roots in this way to remain consistent with our convention that firing moves of two chips at the same site send larger chips to the right and smaller chips to the left. This convention was reversed in [8], which resulted in chips being sorted with the smallest chips on the right and largest chips on the left.

6.2. **Unlabeled Central-Firing.** Just as in labeled chip-firing, every central-firing process has an underlying unlabeled process. Given a central-firing process R on root system Φ, define the corresponding **unlabeled central-firing process** R' by taking every configuration of R modulo the Weyl group of Φ.

Given a configuration λ, taking λ modulo the Weyl group of Φ yields:

- the multiset of chip locations when Φ is of type A_{n-1}.
- the multiset of the absolute values of the chip locations when Φ is of type B_n or C_n.
- the multiset of the absolute values of the chip locations, along with the sign of the product of the chip locations, when Φ is of type D_n.

Note that when Φ is of type A_{n-1}, taking a labeled configuration modulo the Weyl group yields the unlabeled configuration in traditional labeled chip-firing.

We say that a process is **weakly locally confluent** if for any three configurations u, v, and w, such that v and w are reachable from u, there must exist a configuration z that is reachable from v and w. This is weaker than our previously discussed definition of local confluence in that none of these configurations are required to be reachable in a single move. If the process is finite, weak local confluence still does imply global confluence, but the process is no longer required to take a fixed number of moves to complete.

Galashin et al. show that any central-firing process modulo the Weyl group of Φ is weakly locally confluent. In particular, unlabeled central-firing for $\Phi = A_n$ or $\Phi = D_n$ is locally confluent. Unlabeled central-firing for $\Phi = A_n$ is equivalent to chip-firing on the 1-dimensional grid graph, which is known to be locally confluent. Unlabeled central-firing for $\Phi = D_n$ is equivalent to chip-firing on one of the following two graphs:

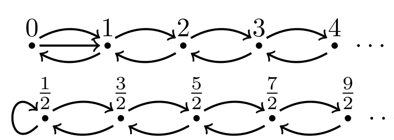

FIGURE 8. For initial weight 0, and ω_1 through ω_{n-2}, unlabeled central-firing on $\Phi = D_n$ corresponds to chip-firing on the top graph. For initial weights ω_{n-1} and ω_n, unlabeled central-firing on D_n corresponds to chip-firing on the bottom graph. For initial weight ω_n, unlabeled central-firing on D_n and B_n are the same, so unlabeled central-firing on B_n for initial weight ω_n also corresponds to chip-firing on the bottom graph.

Because each central-firing process is equivalent to a chip-firing process, central-firing processes with $\Phi = A_n$ or D_n fit into our LLD lattice framework. Given a central-firing process of type A_n or D_n, we first create an **unlabeled configuration poset** by taking the central-firing configuration poset and identifying configurations that are congruent modulo the Weyl group. We can L-color the unlabeled configuration poset as follows:

- For type A_{n-1}, and each site k, let one color consist of all edges corresponding to moves at site k.
- For type D_n, and each $k \geq 0$, let one color consist of all edges corresponding to moves at sites $\pm k$

Consider a site k. By Lemma 4.16 and Corollary 4.17, the colors in the S-coloring corresponding to firing moves at site k (for type A_{n-1}) or at site $\pm k$ (for type D_n) form a chain in the color poset. Thus, we can interpret each color in the S-coloring as "the j^{th} move at site k (or at sites $\pm k$)" for some fixed j.

6.3. **Sorting Conjectures.** Galashin et al. outline all initial weights conjectured to sort for each type of root system central-firing. Their confluence conjecture is as follows:

Conjecture 6.1 ([8], Conjecture 7.1)**.** *Let Φ be of Type A, B, C, or D, and let $\omega \in \Omega \cup \{0\}$ be a fundamental weight or zero. Then central-firing is confluent from ω if and only if $\omega \notin Q + \rho$, unless one of the four exceptional cases happens:*

(1) $\Phi = A_n$, in which case central-firing is confluent from ω if and only if

$$\begin{cases} \omega = 0, \omega_1, \omega_n & n \text{ odd} \\ \omega = \omega_{n/2}, \omega_{n/2+1} & n \text{ even} \end{cases}$$

(2) $\Phi = B_n$, in which case central-firing is confluent from $\omega = \omega_n$ despite the fact that $\omega_n \in Q + \rho$.

(3) $\Phi = D_{4n+2}$ for $n \geq 1$ in which case central-firing is not confluent from $\omega = 0$ even though $0 \notin Q + \rho$.

They prove several results from Conjecture 6.1, but 6 problems are left open.

Conjecture 6.2 ([8], Problem 7.14)**.**

(1) Central-firing is confluent from ω_1 and ω_{2n+1} for $\Phi = A_{2n+1}$.

(2) Central-firing is confluent from ω_n and ω_{n+1} for $\Phi = A_{2n}$.

(3) Central-firing is confluent from ω_n for $\Phi = B_n$ (equivalently, for $\Phi = D_n$).

(4) Central-firing is confluent from $\omega \in \Omega \cup \{0\}$ if and only if $\omega \not\equiv \rho$ in P/Q for $\Phi = C_n$.

(5) Central-firing is not confluent from 0 for $\Phi = D_{4n+2}$.

(6) Central-firing is confluent from $\omega \in \Omega \cup \{0\}$ for all $\omega \not\equiv \rho$ in P/Q for $\Phi = D_n$, except for the case (5) above.

Here, ρ is the Weyl vector of the root system, P is the weight lattice, and Q is the root lattice. We refer the reader to [8] for more details.

The authors have previously made several contributions toward resolving Conjecture 6.2, and we make additional contributions here. The status of Conjecture 6.2 after this paper is as follows:

(1) The proof is summarized by the authors in [11], and proved in full in [12]. The proof is very similar to the proof of (2), as well as to the authors' proof that traditional labeled chip-firing sorts from 0 on A_{2n-1}.

(2) The proof is summarized by the authors in [11], and proved in full in [12].

(3) This is proved in section 6.4 of this paper using LLD lattice-based methods.

(4) This remains open, but our methods are used to provide insight toward this conjecture in [12].

(5) This is not true. We discuss a counterexample and provide an alternative conjecture at the end of this section.

(6) This remains open, but our methods are used to provide insight toward this conjecture in [12].

We now restate Conjecture 6.2 in terms of chips. A labeled chip configuration is **weakly sorted** if for any chips a and b with $a < b$, the position of chip a is less than or equal to the position of chip b. In terms of chips, the open problems are as follows:

(1) Central-firing with $\Phi = A_n$ is confluent from a weakly sorted configuration with 1 chip at the origin and $2n-1$ chips at site 1, or with $2n-1$ chips at the origin and 1 chip at site 1.

(2) Central-firing with $\Phi = A_n$ is confluent from a weakly sorted configuration with n chips at the origin and $n+1$ chips at site 1, or $n+1$ chips at the origin and n chips at site 1.

(3) Central-firing with $\Phi = B_n$ or D_n is confluent from a configuration with all chips at site $\frac{1}{2}$.
(4) Central-firing with $\Phi = C_n$ is confluent from any weakly sorted configuration with an even number of chips at site 1, and the rest at the origin, if the number of chips $n \equiv 1$ or 2 mod 4, or an odd number of chips at site 1, and the rest at the origin, if $n \equiv 0$ or 3 mod 4.
(5) Central-firing with $\Phi = D_n$ is not confluent from any configuration with n chips at the origin, for $n \equiv 2$ mod 4.
(6) Central-firing with $\Phi = D_n$ is confluent from any weakly sorted configuration with an even number of chips at site 1, and the rest at the origin, if the number of chips $n \equiv 2$ or 3 mod 4, or an odd number of chips at site 1, and the rest at the origin, if $n \equiv 0$ or 1 mod 4, with the exception of case (5).

Here, we will use our methods to prove (3). The proof of (3) is similar to our proofs of (1) and (2), as well as to our proof from Section 5 of sorting in traditional labeled chip-firing. In [12], our methods are used to provide insight toward (4), (6), and our modified version of (5).

6.4. **Type B.** We will now prove part (3) of Conjecture 6.2. In initial weight ω_n for $\Phi = B_n$ and D_n, all chips start at site $\frac{1}{2}$. Define a move at location k for $k \geq 0$ to be a firing move involving chips at sites $\pm(k + \frac{1}{2})$. Thus, a type (a) move that sends two chips from $k + \frac{1}{2}$ to $k - \frac{1}{2}$ and $k + \frac{3}{2}$, and a type (d) move that sends chips from sites $\pm(k+\frac{1}{2})$ to sites $k + \frac{3}{2}$ and $-k + \frac{1}{2}$ are both considered to be moves at location k. Note that both types of moves have the same effect on the absolute values of the two chips. If $k \geq 1$, then both send chips from absolute value $k + \frac{1}{2}$ to absolute value $k - \frac{1}{2}$ and $k + \frac{3}{2}$. If $k = 0$, then both send chips from absolute value $\frac{1}{2}$ to absolute value $\frac{1}{2}$ and $\frac{3}{2}$. Note that no type (b) moves may occur in this process, because no chips ever reach the origin throughout the process.

The unlabeled version of this problem corresponds to chip-firing on the bottom graph shown in Figure 8. There is one vertex corresponding to each positive half-integer. From each $k > 0$, there is an edge from $k + \frac{1}{2}$ to $k - \frac{1}{2}$, and an edge from $k - \frac{1}{2}$ to $k + \frac{1}{2}$. There is also a self-loop at $\frac{1}{2}$.

We first need the final unlabeled configuration of this process.

Lemma 6.3. *For every i from 0 to $n-1$, there is exactly one chip in the final configuration whose position has absolute value $i + \frac{1}{2}$.*

Proof. We modify the proof of Theorem 2 from [11]. As shown in [8], the chip configuration modulo the Weyl group of B_n is locally confluent, and thus globally confluent. This implies that the absolute values of the chip positions in any final chip-configuration must be fixed, regardless of the sequence of moves performed. Thus, if we can show a single sequence of moves leading to our desired final configuration, then the same must be true for all possible sequences of moves. We will deal with a sequence that leads to the following intermediate configurations:

$$
\begin{array}{ccccccc}
\frac{1}{2} & \frac{3}{2} & \frac{5}{2} & \frac{7}{2} & \frac{9}{2} & & \\
n & & & & & & \\
n-1 & 1 & & & & & \\
n-2 & 1 & 1 & & & & \\
n-3 & 1 & 1 & 1 & & & \\
n-4 & 1 & 1 & 1 & 1 & & \\
\vdots & & & & & & \\
1 & 1 & 1 & 1 & 1 & \cdots &
\end{array}
$$

In each step, we remove a chip with absolute value $\frac{1}{2}$ and replace it with a chip at the next closest value that is currently empty. Each step only requires the use of 2 chips at $\frac{1}{2}$, so we treat the problem as if there are only 2 chips present there initially. All possible firing moves are made simultaneously in the following pattern:

$\frac{1}{2}$	$\frac{3}{2}$	$\frac{5}{2}$	$\frac{7}{2}$	$\frac{9}{2}$
2	1	1	1	
1	2	1	1	
2	0	2	1	
1	2	0	2	
2	0	2	0	1
1	2	0	1	1
2	0	1	1	1
1	1	1	1	1

There is a line of alternating 2's and 0's (or 1's at site $\frac{1}{2}$) that expands until an extra chip is deposited at the farthest point, and then contracts again. Starting with n chips and then repeating the above process until there is only one chip left at $\frac{1}{2}$ yields a configuration of the desired form.

□

Now, define move k_j to be the j^{th} to last move at location k, where a move at location k is again defined to be any move involving chips moving from sites $\pm(k+\frac{1}{2})$. We want to prove the existence of a similar diamond to what exists in the type A case.

Define P_{ω_n} to be the poset of unlabeled configurations. Here, we mod out by the Weyl group, so that each element in the configuration poset corresponds to the multiset of absolute values of the chip positions. We would like to use join-irreducibles of P_{ω_n} to provide information about the move poset, so we first need to prove that P_{ω_n} is an LLD lattice.

Lemma 6.4. *Consider the unlabeled configuration poset P_{ω_n} of a central-firing process of Type B_n with initial weight ω_n. Color P_{ω_n} such that each color corresponds to the location of each firing move. This coloring is a valid L-coloring.*

Proof. Given an unlabeled configuration u and location k, there is at most one way to perform a firing move at k from configuration u (note that type (a) and (d) firing moves at k have the same effect on u). Thus, this coloring satisfies condition (1) of L-coloring.

Furthermore, as each unlabeled firing move corresponds to a chip-firing move on the bottom graph shown in Figure 8, the process is locally confluent (in the stronger sense discussed in the preliminaries), so the coloring satisfies condition (2), and is thus an L-coloring. □

Lemma 6.5. *Consider the unlabeled configuration poset P_{ω_n} of a central-firing process of Type B_n with initial weight ω_n. Then each color in the S-coloring of P_{ω_n} consists of all edges corresponding to moves of the form "the j^{th} firing move at location k," for fixed location k and integer j.*

Proof. By Lemma 6.4, there exists a valid L-coloring in which each color corresponds to a single location being fired. By Lemma 4.15, each color in this L-coloring is a union of colors of the S-coloring which form a chain in the color poset. As a result, given a location k, all colors of the S-coloring corresponding to firing moves at k must occur in the same order in any complete downward path in P_{ω_n}. Thus, each color in the S-coloring corresponding to location k must consist of all moves of the form "the j^{th} move at k" for some fixed j. □

Becuase the colors of the S-coloring correspond to moves of the move poset, the move poset is isomorphic to the poset of join-irreducibles of P_{ω_n}. Define $J(k_j)$ to be the join-irreducible configuration corresponding to move k_j. The rest of the proof is similar to the proof of sorting for labeled chip-firing with $2m$ chips. We prove the existence of a locally confluent diamond of firing moves, and we show that the diamond forces chips into their final sorted positions.

Lemma 6.6. *For $1 \le j \le n-k-1$, move k_j must take place before $(k+1)_{j-1}$ and $(k-1)_j$, if sites $k+1$ and $k-1$ fire at least $j-1$ times and j times, respectively. Further, move k_j must take place with exactly 2 chips present at sites $\pm(k+\frac{1}{2})$.*

Proof. The proof closely follows Lemmas 5.4 and 5.5. We identify the join-irreducible configurations associated with each of the moves in question, and we can show reachability relations between the relevant join-irreducibles.

By Theorem 4.12, a move $(k_1)_{j_1}$ being required to take place before $(k_2)_{j_2}$ is equivalent to $J\big((k_2)_{j_2}\big)$ being reachable from $J\big((k_1)_{j_1}\big)$. As a result, we will show that for $1 \le j \le n-k$, join-irreducibles $J\big((k+1)_{j-1}\big)$ and $J\big((k-1)_j\big)$ are reachable from $J(k_j)$, if those elements exist.

We will show that for $k \ge 0$, $0 \le j \le n-k-1$, $J(k_j)$ consists of 2 chips at location k, 0 chips at location $k+j$, and 1 chip at all other locations from 0 to $n-1$.

A configuration with the specified numbers of chips at each location is reachable through the following firing moves at each location:

- all moves at locations $k+j$ through $n-1$
- all but the last j moves at locations 0 through k
- all but the last $k+j-x$ moves at locations $x = k+1$ through $x = k+j-1$

We can show that these firing moves produce the desired configuration by comparing the number of firing moves at each location and its neighbors to the total number of firing moves throughout the process.

- At locations $x > k+j$, all moves have occurred at location x and its neighbors, so the number of chips at location x in configuration $J(k_j)$ is equal to the number of chips (1 chip) at location x in the final configuration.
- At location $x = k+j$, all moves have occurred at location x and one of its neighbors, but one move has not yet occurred at its other neighbor. As a result, location x has one fewer chip in configuration $J(k_j)$ (0 chips) than location x has in the final configuration.
- The number of firing moves yet to occur at location k exceeds the average number of remaining firing moves at its neighbors by $\frac{1}{2}$, so location k has one more chip in configuration $J(k_j)$ (2 chips) than location k has in the final configuration.
- At all other locations x, the number of firing moves yet to occur at location x is equal to the average number of remaining firing moves at each neighbor of location x, so location x has as many chips in configuration $J(k_j)$ (1 chip) as location x has in the final configuration.

As a result, each join-irreducible takes the desired form. Now consider join-irreducibles k_j and $(k+1)_{j-1}$, for $k > 0$. $J(k_j)$ consists of 2 chips at location k, 0 chips at location $k+j$, and 1 chip at all other locations from 0 to $n-1$. $J\big((k+1)_{j-1}\big)$ consists of 2 chips at location $k+1$, 0 chips at location $k+j$, and 1 chip at all other locations from 0 to $n-1$. It is possible to get from $J(k_j)$ to $J\big((k+1)_{j-1}\big)$ by firing successively at every location from k down to 0, so $J\big((k+1)_{j-1}\big)$ is reachable from $J(k_j)$. Similarly, it is possible to get from $J(k_j)$ to $J\big((k-1)_j\big)$ by firing successively at every location from k up to $k+j-1$, so $J\big((k-1)_j\big)$ is reachable from $J(k_j)$.

Because the corresponding join-irreducibles satisfy the desired reachability constraints, the moves of the poset satisfy the corresponding order, as desired.

For each k_j considered, $J(k_j)$ has 2 chips at location k. Because all firing moves at neighboring locations that can occur before $J(k_j)$ must do so to reach configuration $J(k_j)$, the number of chips at location k in configuration $J(k_j)$ must be the maximum number of chips that can be available for move k_j. Since move k_j cannot take place with fewer than 2 chips at location k, it must take place with exactly 2 chips present. $\square$

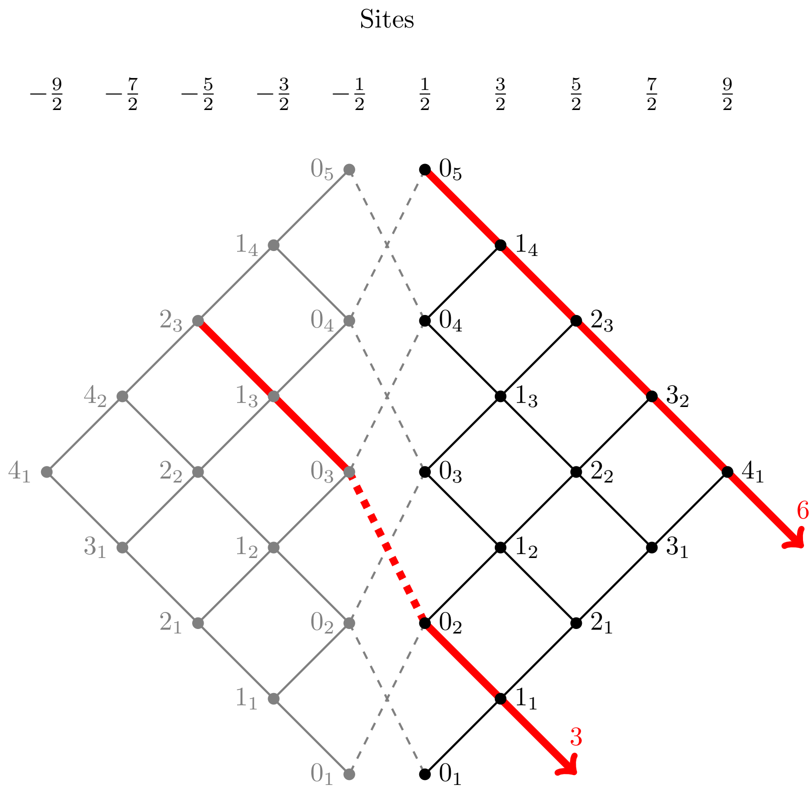

Figure 9. The Hasse diagram of the firing moves at the end of the labeled chip-firing process on B_6 (in which all chips begin at site $\frac{1}{2}$), and consequently, the join-irreducibles of the unlabeled configuration poset. The moves are duplicated in gray on the left side of the figure to represent that some moves may involve a chip initially at a negative site. The "paths" followed by chips 3 and 6 are shown in red. Neither chip may be involved in a firing move to the left of its corresponding red line, and must be fired to the right any time it is involved in a firing move on the red line.

Now that we have results on the unlabeled system, number the chips from 1 to n. We can establish lower bounds on the positions of various chips throughout the process.

Lemma 6.7. *For any k from 1 to n, the position of chip k must never be less than $\frac{1}{2}+k-n$ at any point in the chip-firing process.*

Proof. Proceed by strong induction on k, going downward from n. $k=n$ is the largest possible label that a chip can have. Thus, it can never be fired to the left from its initial position at $\frac{1}{2}$.

Now suppose that for each chip k' with label greater than k, the position of k' may never be less than $\frac{3}{2}+k-n$. Then consider the chip with label k. If the chip ever reaches the position $\frac{1}{2}+k-n$, then by our inductive assumption, it must be the smallest chip at that site. Thus, it cannot be sent further to the left from that site, and its position can never be less than $\frac{1}{2}+k-n$ as desired. The result follows by induction. $\square$

The next lemma tracks the locations of chips as they "move through the diamond." Given a chip's position throughout the process, we can use the relationships between the various firing moves to guarantee that specific chips must be sufficiently far to the right at specific times.

Lemma 6.8. *For each chip k with $2 \leq k \leq n$, and each x with $0 \leq x \leq n-1$, the position of chip k must be at least $\frac{1}{2}+k-n+x$ immediately preceding firing move $(n-k-x-1)_{k-1}$ for $x < n-k$, or $(k-n+x)_{n-x}$ for $x \geq n-k$.*

Proof. Proceed by induction, first on x, and then on k. First, let $k = n$ and $x = 0$. By Lemma 6.7, the position of chip n may never be less than $\frac{1}{2}$. In particular, chip n must be at position at least $\frac{1}{2}$ prior to move 0_n. Now suppose that chip n is at position at least $\frac{1}{2} + x$ immediately preceding move x_{n-x}. If chip n is at position $\frac{1}{2} + x$, then because there can only be two chips present at sites with absolute value $\frac{1}{2} + x$ prior to move x_{n-x}, chip n must be fired by move x_{n-x}, and since it is the largest chip at site $\frac{1}{2}+x$, it must be fired to the right. Thus, chip n must be at position at least $\frac{1}{2} + x + 1$ immediately following move x_{n-x}. Since move x_{n-x} must occur between move $(x+1)_{n-x}$ and move $(x+1)_{n-x-1}$, chip n must be at position at least $\frac{1}{2} + x + 1$ by the time the move $(x + 1)_{n-x-1}$ occurs. Thus, the bounds on chip n follow by induction.

Now, induct on k. Assume that the position bounds hold for all chips $k' > k$ and $0 \leq x \leq n - 1$. We show that the bounds hold for chip k and all values for x from 0 to $n - 1$. Chip k can never reach a position less than $\frac{1}{2} + k - n$ at any point in the firing process, and in particular, chip k must be at position at least $\frac{1}{2} + k - n$ immediately preceding move $(n - k - 1)_{k-1}$. Thus, the bound holds for $x = 0$. Now, assume that chip k is at position at least $\frac{1}{2}+k-n+x$ immediately preceding firing move $(n-k-x-1)_{k-1}$ for $x < n-k$. By our inductive assumption, all chips with value greater than k must already be at positions greater than $\frac{1}{2} + k - n + x$ when move $(n - k - x - 1)_{k-1}$ occurs. Thus, if chip k is at position $\frac{1}{2} + k - n + x$, then it must be the greatest chip at that position, and because there are at most two chips present for firing move $(n-k-x-1)_{k-1}$, chip k must be fired to the right. Chip k thus must be at position at least $\frac{1}{2}+k-n+x+1$ by the time the next move occurs at that site. If $x < n - k - 1$, the next move at site $(n - k - x)$ is move $(n - k - x)_{k-1}$. If $x = n - k - 1$, then the next move at site $(n - k - x)$ is 0_k. The same proof applies for $x \geq n - k$, so these position bounds hold for all $2 \leq k \leq n$ and all $0 \leq x \leq n - 1$, as desired. ☐

Lemma 6.8 allows us to complete our final result.

Theorem 6.9 ([8], Problem 7.14, Part (3))**.** *When the labeled chip-firing process terminates on a system of type B_n (or D_n), with initial configuration ω_n, each chip k, $2 \leq k \leq n$, is at position $k - \frac{1}{2}$.*

Proof. By Lemma 6.8, chip k must be at position at least $k - \frac{3}{2}$ by the time move $(k - 2)_1$ occurs. Using the same argument as in Lemma 6.8, chip k must be at position at least $k - \frac{1}{2}$ immediately following move $(k - 2)_1$. Since no firing moves may occur at positions greater than $k - \frac{3}{2}$ after firing move $(k - 2)_1$, the final position of chip k must be greater than or equal to $k - \frac{1}{2}$. The only way for every chip to satisfy this condition is if every chip k, $2 \leq k \leq n$ is at position $k - \frac{1}{2}$ in the final configuration, as desired. ☐

Note that chip 1 may be at position $\pm\frac{1}{2}$, depending on the value of n, but as chip 1 is the smallest chip, all chips will be in sorted order regardless.

Theorem 6.9 on labeled chip-firing actually allows us to prove a previously open problem about unlabeled chip-firing. Local confluence modulo the Weyl Group is enough to guarantee that the final chip-configuration consists of 1 chip at $\pm\frac{1}{2}$, 1 chip at $\pm\frac{3}{2}, \ldots$ and 1 chip at $\pm n - \frac{1}{2}$. However, this alone does not guarantee whether those chip locations are positive or negative. Our result, however, does make this guarantee. Other than possibly the chip at $\pm\frac{1}{2}$, all chips must terminate at positive final positions.

Corollary 6.10. *Consider a collection of n indistinguishable chips on the line, beginning with all chips at position $\frac{1}{2}$. Let the system evolve using two types of moves:*

(1) If two chips are at the same site, send one of the chips at that site one unit to the left, and one of the chips at that site one unit to the right.
(2) If two chips are at sites with opposite coordinates, send both chips one unit to the right.

Then the process terminates with 1 chip at every half-integer site from $\frac{3}{2}$ to $n - \frac{1}{2}$. Furthermore, if $n \equiv 0$ or 1 mod 4, then the final position also has a chip at site $-\frac{1}{2}$. If $n \equiv 2$ or 3 mod 4, then the final position also has a chip at site $\frac{1}{2}$.

Proof. Assign labels 1 through n to the chips. Whenever a move of type (2) is performed, any chips at the two firing sites may be chosen to be sent to the right. Whenever a move of type (1) is performed, any two chips at the firing site may be chosen to fire, with the chip with the larger label sent to the right, and the chip with the smaller label sent to the left. This process is identical to root system central-firing of type B_n (or D_n), so by Theorem 6.9, the process terminates with 1 chip at every half-integer site from $\frac{3}{2}$ to $n - \frac{1}{2}$. As the addition of labels does not restrict which chips can move at which times, the original unlabeled process must also terminate with chips in positions $\frac{3}{2}$ to $n - \frac{1}{2}$.

Now, note that every firing move of type (1) does not change the sum of the chip positions, while every firing move of type (2) increases the sum of the chip positions by 2. Thus, the difference between the final sum of the chip positions and the initial sum must be even. By local confluence of unlabeled chip-firing, the last chip must be at either site $-\frac{1}{2}$ or site $\frac{1}{2}$.

The initial sum of the chip-positions is equal to $\frac{n}{2}$, and the final sum is equal to $\frac{n^2}{2}-1$ (if the remaining chip is at $\frac{1}{2}$) or $\frac{n^2}{2}$ (if the remaining chip is at $\frac{1}{2}$). Now, consider 4 cases:

$$\begin{array}{lr}
n \equiv 0 \bmod 4 & \frac{n^2}{2} - \frac{n}{2} = \frac{n}{2}(n-1) \text{ (even)} \\
n \equiv 1 \bmod 4 & \frac{n^2}{2} - \frac{n}{2} = \frac{n-1}{2}(n) \text{ (even)} \\
n \equiv 2 \bmod 4 & \frac{n^2}{2} - 1 - \frac{n}{2} = \frac{n-2}{2}(n+1) \text{ (even)} \\
n \equiv 3 \bmod 4 & \frac{n^2}{2} - 1 - \frac{n}{2} = \frac{n+1}{2}(n-2) \text{ (even)}
\end{array}$$

Note that as $\frac{1}{2}$ and $-\frac{1}{2}$ differ by 1, the difference involving $\frac{1}{2}$ is even if and only if the difference involving $-\frac{1}{2}$ is odd. Thus, the only way to produce an even difference is if the last chip is at $\frac{1}{2}$ ($n \equiv 0$ or 1 mod 4) or at $-\frac{1}{2}$ ($n \equiv 2$ or 3 mod 4). □

Note that while the final location of the smallest chip depends on n, it does not depend on the choices of firing moves throughout the process. Thus, the final chip positions do not depend on the choices of moves.

In addition to resolving a conjecture about labeled chip-firing, we have used a labeled chip-firing result to prove an unlabeled chip-firing result. This makes Corollary 6.10 particularly notable, as it shows the value of labeled chip-firing to the broader field of chip-firing.

Remark 6.11. Methods based on move posets can also be used to study root system central-firing of Types C_n and D_n and provide insight into parts (4) through (6) of Conjecture 6.2.

The authors have shown computationally that part (5) is not true, as central-firing on $\Phi = D_{10}$ is confluent from 0 (as well as the more trivial $\Phi = D_2$). From the specific nature of the way in which central-firing fails to sort from 0 for $\Phi = D_6$, we conjecture instead that central-firing on $\Phi = D_6$ is not confluent from 0, and that this is the only exception to part (6) of Conjecture 6.2.

See [12] for progress on parts (4) and (6), including details about the structures of the move posets in Types C_n and D_n. There is a clear distinction between the move posets of cases conjectured to sort and those conjectured not to sort.

The result of this work leaves us with the following conjecture for the remaining open problems in central-firing.

Conjecture 6.12.
(1) *Central-firing is confluent from $\omega \in \Omega \cup \{0\}$ if and only if $\omega \not\equiv \rho$ in P/Q for $\Phi = C_n$.*
(2) *Central-firing is confluent from $\omega \in \Omega \cup \{0\}$ for all $\omega \not\equiv \rho$ in P/Q for $\Phi = D_n$, except for the case $\Phi = D_6$, $\omega = 0$.*

Part (1) is equivalent to part (4) from the original conjecture of [8]. Part (2) is based on parts (5) and (6) from the original conjecture, although (5) has been modified due to our counterexample for central-firing when $\Phi = D_{10}$.

7. Distributive Lattices

As every distributive lattice is LLD, the results in Section 4 apply to all distributive lattices. Here, we consider some classes of distributive lattices previously studied by Propp [16]. Specifically, we analyze distributive lattices generated by graph orientations, bipartite matchings and tilings, and spanning trees.

7.1. **Graph Orientations.** Given a graph G, an orientation R of the edges of G, and a directed cycle C, define the **circulation** of R around C as the number of edges in R oriented in the same direction as C, minus the number of edges of R oriented in the opposite direction. Define the **circulation** of R to be a function r that maps each cycle C to the circulation $r(C)$ of r around C. We say that R is an r**-orientation**.

Define an **accessibility class** of R as a set of vertices that can be reached from one another through the oriented edges in R. An accessibility class is **maximal** if all directed edges in R between A and A^c point toward A, and it is **minimal** if they all point toward A^c. If A is maximal, then the operation of reversing all edges between A and A^c (so that they now point toward A^c) is called **pushing down** A.

Push-down operations allow us to define a distributive lattice on the set of orientations with a fixed circulation r. Designating a special accessibility class A^*, Propp shows that we can define a partial ordering on the r-orientations of G, in which r-orientation R covers r-orientation S when S can be obtained from R by pushing down an accessibility class other than A^*. Propp proves that the poset generated by this cover relation is a distributive lattice. This poset further defines a process, in which each orientation is a configuration, and each push-down operation is a move.

We show how to interpret the results of Section 4 on this lattice of orientations. First, we show that coloring each move based on which accessibility class is pushed down is a valid L-coloring.

Lemma 7.1. *Given a graph G and a circulation r, then color the poset of r-orientations such that each color corresponds to an accessibility class A being pushed down. Then this coloring is a valid L-coloring.*

Proof. First, given an orientation R and an accessibility class A, there is at most one way to push down A from R. Thus, this coloring satisfies condition (1) of L-coloring.

Now, consider an orientation R and two accessibility classes A_1 and A_2 that can be pushed down from R. If both can be pushed down, then all directed edges between A_1 and A_1^c are directed toward A_1, and the same is true for A_2. Thus, there can be no edges between A_1 and A_2, so pushing one down will not affect the other. As a result, the push-down operations can be performed in either order to reach the same resulting configuration, so the coloring satisfies condition (2), and is thus an L-coloring. □

Theorem 7.2. *Given a graph G and circulation r, consider the distributive lattice of r-orientations on G. Then each color in the S-coloring of this lattice consists of all edges corresponding to moves of the form "the j^{th} push-down operation at class A" for a fixed accessibility class A and positive integer j.*

Proof. By Lemma 7.1, there exists a valid L-coloring in which each color corresponds to a single accessibility class being pushed down. By Lemma 4.15, each color in this L-coloring is a union of colors of the S-coloring which form a chain in the color poset. As a result, given an accessibility class A, all colors of the S-coloring corresponding to push-down operations at A must occur in the same order in any complete downward path in the lattice of r-orientations. Thus, each color in the S-coloring corresponding to class A must consist of all moves of the form "the j^{th} move at A" for some fixed j. □

7.2. **Tilings.** Propp shows that bipartite matchings on planar graphs also form distributive lattices by relating them to processes involving r-orientations [16].

Given a planar bipartite graph $G = (V, E)$ on the sphere, and a function $d : V \to \mathbb{N}$, define a d**-factor** of G as a subgraph of G (viewed as a set of edges) in which each vertex v has degree $d(v)$. The example we are most concerned with is a 1-factor, in which $d(v) = 1$ for all $v \in G$, corresponding to a perfect matching of G.

Color the vertices of one part of the bipartite graph white, and the vertices of the other part black. Define an **elementary cycle** as a cycle encircling a single face of the planar graph. Given a d-factor M, an **alternating cycle** in G relative to M is an elementary cycle in G in which the edges alternatively belong to M and M^c. Call the cycle **positive** if the edges in M, oriented clockwise around the cycle, go from black vertices to white vertices, and negative if they go from white vertices to black vertices. Given a d-factor M, we can perform a **twisting down** operation by taking a positive alternating cycle, and replacing the edges of M in that cycle with the edges of M^c in that cycle, thus turning the cycle from a positive to a negative alternating cycle. Define **twisting up** to be the inverse of twisting down.

Designate a special face f^* of G. Since G is on a sphere, we usually let f^* be the "outside" face when G is viewed on a plane. Propp shows that the set of d-factors can be partially ordered such that a d-factor M_1 covers a d-factor M_2 if it is possible to get from M_1 to M_2 by twisting down a face other than f^*. He further shows that this partial ordering forms a distributive lattice. We can show that coloring the edges of this lattice according to which face is twisted down, forms an L-coloring.

Lemma 7.3. *Given a bipartite planar graph G and an $\mathbb{N}$-valued function d on the vertices, along with special face f^*, consider the lattice of d-factors of G with special face f^*. Then the edge coloring of the lattice such that each color corresponds to the face being twisted down, is a valid L-coloring.*

Proof. First, given a d-factor M and a face f, there is at most one way to twist down f from M. Thus, this coloring satisfies condition (1) of L-coloring.

Now, consider a d-factor M and two faces f_1 and f_2 that can be twisted down from M. If both can be twisted down down, then all black-white edges in f_1 and f_2 must be oriented clockwise. Since an edge is oriented clockwise in f_1 if and only if it is oriented counterclockwise in the face that borders f_1 along that edge, f_1 and f_2 must not be adjacent. Thus, twisting down f_1 does not affect whether f_2 can be twisted down and vice versa, so the operations can be performed in either order. As a result, this coloring satisfies condition (2), and thus is an L-coloring. □

We now proof that the moves of the S-coloring of a d-factor lattice correspond to move types of the form "j^{th} twisting down operation at face f." This proof is nearly identical to the proof of Theorem 7.2.

Theorem 7.4. *Given a bipartite planar graph G and and an $\mathbb{N}$-valued function d on the vertices, along with special face f^*, consider the distributive lattice of d-factors on G with respect to f^*. Then each color in the S-coloring of the lattice consists of all edges corresponding to moves of the form "the j^{th} twist-down operation at face f" for a fixed face f and positive integer j.*

Proof. By Lemma 7.3, there exists a valid L-coloring in which each color corresponds to a single face being twisted down. By Lemma 4.15, each color in this L-coloring is a union of colors of the S-coloring which form a chain in the color poset. As a result, given a face f, all colors of the S-coloring corresponding to twist-down operations at f must occur in the same order in any complete downward path in the lattice of d-factors. Thus, each color in the S-coloring corresponding to class face f must consist of all moves of the form "the j^{th} move at f" for some fixed j. □

As a result of Theorem 7.4, the distributive lattices generated by bipartite matchings (and as a consequence, domino and lozenge tilings of simply connected regions in 2 dimensions) fall under our framework. The framework provides us with some nice results on paths between tilings. In particular, Theorem 4.19 and Lemma 4.21 provide a method for finding shortest paths between tilings using flip moves, as well as showing which flip moves must be performed on a shortest path between any two flip-connected tilings. Previous approaches to this result considered the graph of flip-connected tilings as a CW-complex [17], and then showed that the complex was contractible. Here, the distributive lattice structure allows us to show these results on shortest paths using colorings.

Given a bipartite planar graph G with function d on the vertices and special face f, consider the distributive lattice generated by the twist down relation, and its maximum element $\mathbf{1}$. Suppose that the S-coloring of the lattice has colors $c^1, c^2, \dots c^n$. Define the color vector of d-factor u_1 as $V(u_1) = (v_1(u_1), v_2(u_1), \dots v_n(u_1))$ such that for each i from 1 to n, v_i is 1 if any shortest path from $\mathbf{1}$ to u_1 contains color c^i, and 0 otherwise. Similarly define a color vector of a path s as $V(s) = (v_1(s), v_2(s), \dots v_n(s))$ such that for each i from 1 to n, $v_i(s)$ is the number of downward edges of color c^i in s, minus the number of upward edges of color c^i in s. The following results follow directly from their counterparts in Section 4:

Theorem 7.5. *Consider a path s from tiling u_1 to tiling u_2 on a simply connected region G. Then $V(s) = V(u_2) - V(u_1)$. Thus, all paths from u_1 to u_2 contain the same flip moves, up to moves being done and then undone.*

Theorem 7.5 allows us to uniquely define a color vector $V(u_1, u_2)$, equal to the color vector of any path s from u_1 to u_2.

Corollary 7.6. *For any two tilings u_1 and u_2 on simply connected region G, all entries of $V(u_1, u_2)$ are in $\{-1, 0, 1\}$.*

Lemma 7.7. *Given tilings u_1 and u_2 on simply connected region G, there exists a path from u_1 to u_2 with length $|V(u_1, u_2)|_{L^1}$. This is the shortest possible length of any path from u_1 to u_2. Furthermore, there exists a path of length $|V(u_1, u_2)|_{L^1}$ from u_1 to u_2 in which all downward edges precede all upward edges, and there also exists a path of length $|V(u_1, u_2)|_{L^1}$ from u_1 to u_2 in which all upward edges precede all downward edges.*

Proof. This result follows from Lemma 4.21. The fact that a shortest path exists with upward moves followed by downward moves is due to the fact that the dual of a distributive lattice is still distributive, and hence LLD. □

We conclude this section by providing a lattice-based proof of the following result of Saldanha et al. (equivalent to Theorem 3.4 of [17]).

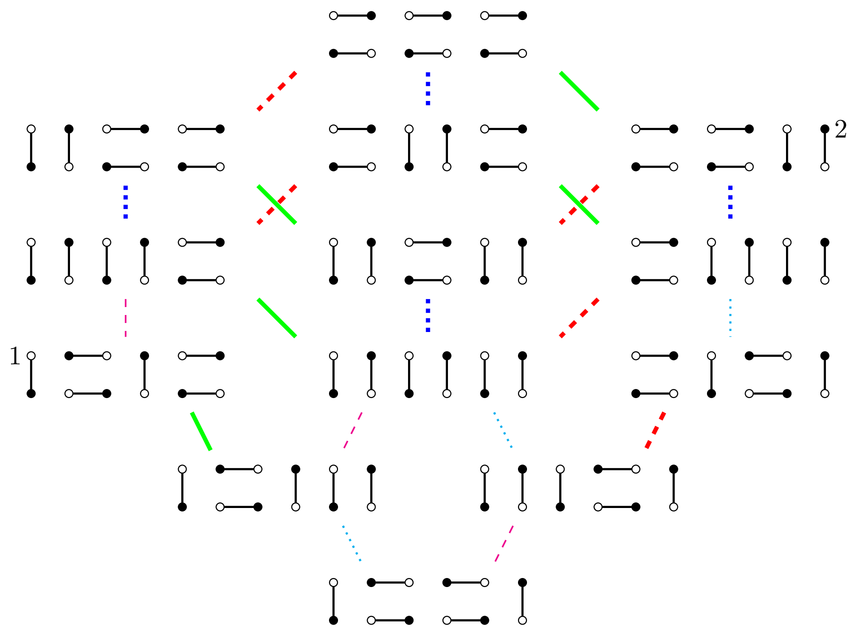

FIGURE 10. The partial ordering of bipartite matchings (or equivalently, domino tilings) on a 2x6 grid. As each face is only twisted down once to get from the top to the bottom, each edge color corresponds to twisting down a different face. The distance between the configurations labeled 1 and 2 in this poset is 4, and each shortest path from 1 to 2 must contain upward edges colored red, blue, and magenta, and a downward edge colored green.

Theorem 7.8. *Consider tilings u_1 and u_2 on simply connected region G, and two shortest paths s_0 and s_f from u_1 to u_2. Then it is possible to get from s_0 to s_f through a sequence of edge pair swaps.*

Proof. Suppose that the tilings of G have configuration poset P with S-coloring c. We first show that it is possible to get from s_0 to some path s_0' in which all downward edges precede all upward edges. Let $s_0 = (e^1, e^2, \dots e^n)$. Suppose that there exists some index i such that e_i is a downward edge, but e_{i+1} is an upward edge. By property (2) of L coloring, there must be a downward edge of color $c(e_{i+1})$ followed by an upward edge of color $c(e_i)$ from the start point of edge e_i. Thus, we can perform an edge pair swap to reverse the order of colors $c(e_i)$ and $c(e_{i+1})$, resulting in a downward edge moving before an upward edge in the sequence, without changing the numbers of downward and upward edges. This operation reduces the sum of the indices of the downward edges by 1. Since that sum can never be lower than $\binom{\text{number of downward edges }+1}{2}$, this process must terminate with a path s_0' in which all downward edges precede all upward edges. Similarly, it is possible to get from s_f to a path s_f' in which all downward edges precede all upward edges.

By Lemma 4.19, the downward edges (and the upward edges respectively) of s_0' and s_f' must be the same. By Theorem 4.12, each configuration has a single, unique color vector, so s_0' and s_f' must follow i downward edges from u to a common configuration w, and then follow $n - i$ upward edges to v. By Lemma 4.14, it is possible to transform the first i edges in s_0' to the first i edges in s_f' via edge pair swaps, and it is possible to transform the last $n - i$ edges in s_0' to the last $n - i$ edges in s_f' via edge pair swaps.

Thus, it is possible to transform from s_0 to s_0' to s_f' to s_f via edge pair swaps, as desired. □

A domino tiling example on a 2x6 grid is shown in Figure 10.

7.2.1. *Tilings as Chip-Firing.* One interesting consequence of Propp's correspondence between graph orientations and bipartite matchings is that it also provides us with a correspondence between chip-firing and bipartite matchings.

Propp [16] provides a method for converting any connected bipartite planar graph G into an orientation of its dual graph $G^\perp$. To do this conversion, color all of the vertices in one part of the bipartite graph black, and color all of the vertices in the other part white. Define the standard orientation on $G^\perp$ as the orientation

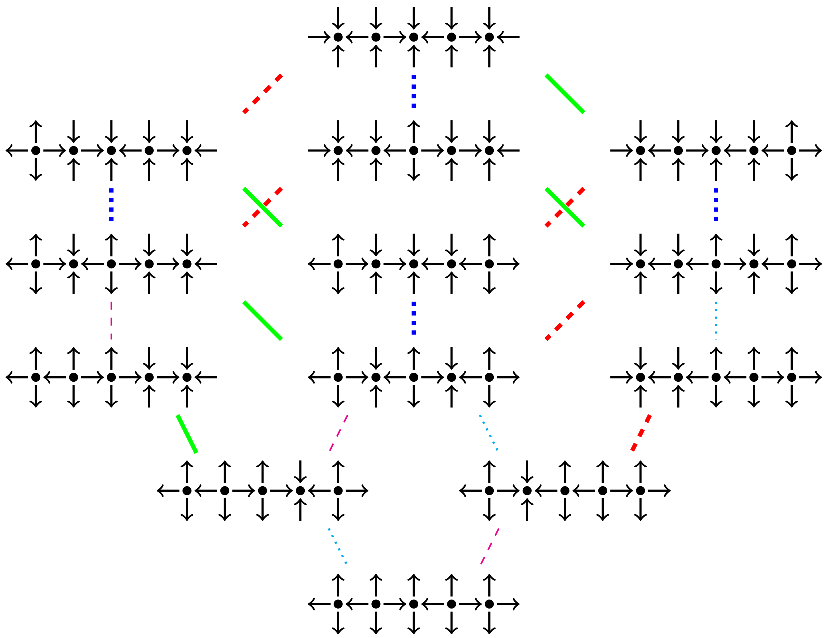

FIGURE 11. The graph orientations corresponding to each bipartite matching in Figure 10. Edges without a second endpoint represent edges in which one endpoint is the "outside" of $G^\perp$

in which each edge e' in $G^\perp$ corresponding to edge e in G is oriented so that the white endpoint of e is to its left and the black endpoint of e is to its right. Given a d-factor u on G, create a corresponding orientation on $G^\perp$ such that all edges in the standard orientation corresponding to edges in u are reversed, and all edges in the standard orientation not corresponding to edges in u have their directions maintained. Propp shows that a pushing down operation in the dual orientation corresponds to a twisting down operation in G, and that the pushing down operations generate the same distributive lattices as the twisting down operations. The set of d factors corresponding to a particular function d corresponds to the set of orientations with a particular circulation on $G^\perp$.

Now, any graph orientation R can be converted to a chip-configuration by assigning to each vertex a number of chips equal to its number of incoming edges in R. As noted in Theorem 3.3 of [4], a vertex can fire if and only if all of its incident edges in R are directed inward, in which case firing corresponds to reversing the directions of all incident edges in R, or, in the language of Propp, pushing down that vertex.

We further note that pushing down any accessibility class A corresponds to cluster-firing all of the vertices in that class in the same manner. A cluster fire of cluster A can occur if and only if each vertex v in A has at least as many chips as it has edges to A^c, in which case a cluster fire reduces the number of chips at v by the number of edges between v and A^c, sending one chip along each edge to A^c. Cluster firing A corresponds to reversing the direction of every edge between A and A^c in the corresponding orientation.

Thus, we can create a correspondence between d-factors (or, more specifically, bipartite matchings or domino tilings) and chip-firing processes. Take a connected bipartite planar graph G with a distinguished region f^*, which corresponds to a sink vertex in the corresponding chip-firing process. Then use Propp's methods to convert this d-factor to an orientation of the graph $G^\perp$, and then convert from that graph orientation to a chip configuration.

Each push-down operation at a face f in G corresponds to a firing move at the corresponding vertex in $G^\perp$. In particular, flip moves in a domino tiling correspond to firing moves (or un-firing moves) in the corresponding chip-firing process. This implies that the distributive lattice generated by push-down operations is the same as the lattice generated by a chip-firing process with a particular initial configuration on the dual graph. The graph orientations and chip configurations generated by the domino tiling system in Figure 10 are shown in Figure 11 and Figure 12.

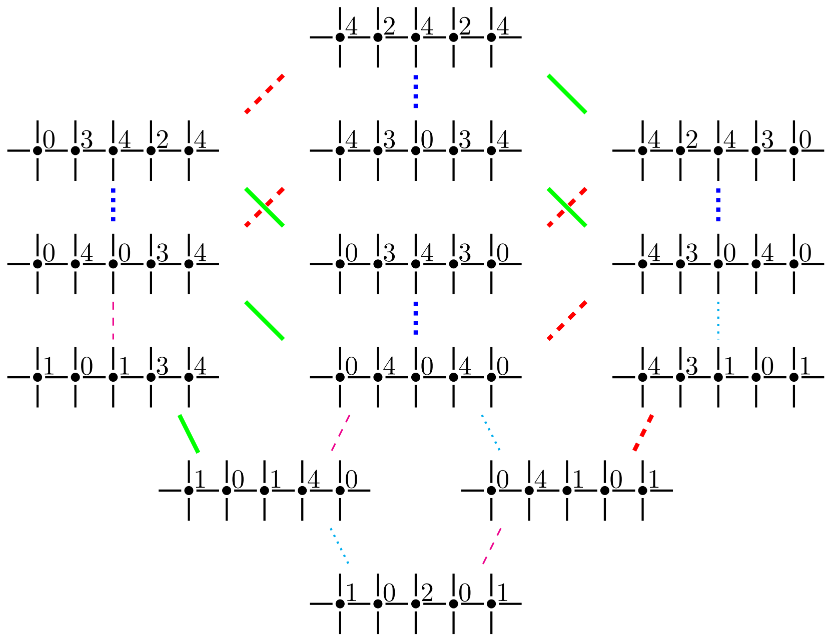

FIGURE 12. The chip configurations corresponding to the graph orientations in Figure 10. Edges without a second endpoint correspond to edges to the sink.

7.3. **Spanning Trees.** Propp also defines a distributive lattice on the set of spanning trees of a planar graph. Trees are related by a successor relation involving a "swinging down" move that transforms one spanning tree to another through rotating one edge clockwise about a vertex, while meeting certain other conditions, as described below.

Given a planar graph G, designate a special face f^*, with a special vertex v^* adjacent to f^*. Then consider a vertex v and an edge e containing v. Define the **clockwise successor** e' of e with respect to v to be the next edge coming out of v if we move clockwise from e through face f. Given v, e, f, and clockwise successor e', and spanning tree T, angle eve' is **positively pivotal** if all of the following are met:

(1) $e \in T$ and $e' \notin T$.
(2) $T' = T \backslash \{e\} \cup \{e'\}$ is a spanning tree of G.
(3) The simple path from v to v^* in T contains e.
(4) The simple path from v to v^* in T' contains e'.
(5) The (unique) simple cycle in $T \cup T'$ separates f from f^*.

If eve' is positively pivotal, then the operation transforming T to T' is called **swinging down** at v through angle eve'. Transforming T' to T is called swinging up. Propp shows that every choice of f^* and v^* generates a partial ordering of spanning trees on G with a cover relation $T \gtrdot T'$ if it is possible to get from T to T' through one swing-down operation. He further shows that the poset generated by this cover relation is a distributive lattice. The proof that this poset is a distributive lattice involves showing that a swinging down operation through a specific angle eve' corresponds to a twisting down operation along a specific face of a graph $H(G)$. See [16] for more details.

We can show that coloring the edges of this lattice of spanning trees according to the angle of the corresponding swinging down operation, is an L-coloring.

Lemma 7.9. *Given a planar graph G with special face f^* and special vertex v^* on f^*, consider the distributive lattice P of spanning trees generated by swinging down operations. Let c be an edge coloring of P such that each color of c includes all edges that correspond to swinging down through a particular angle $A = eve'$. Then c is a valid L-coloring.*

Proof. Let $H(G)$ be the Hasse diagram of the face-poset of G, viewed as a graph. $H(G)$ is a graph with a vertex $\bar{v}$ corresponding to each vertex v of G, a vertex $\bar{e}$ corresponding to each edge e of G, and a vertex $\bar{f}$ corresponding to each face f of G. For each edge e of G, and each endpoint v of e, there is an edge in $H(G)$

between $\bar{e}$ and $\bar{v}$. For each face f of G and each edge e adjacent to f, there is an edge in $H(G)$ between $\bar{f}$ and $\bar{e}$. This idea was introduced by Temperly [19], discussed by Lovasz [13], and generalized by Burton and Pemantle [6] to infinite planar graphs. Propp shows that the lattice of spanning trees of G is isomorphic to the lattice of matchings of $H(G)$, such that swinging down through an angle eve' in G corresponds to twisting down a unique face of $H(G)$ [16].

As swinging down through eve' corresponds to twisting down a unique face of $H(G)$, coloring P according to the swinging angle is equivalent to coloring a lattice of d-factors on $H(G)$ with a coloring c', in which each color corresponds to the face of $H(G)$ that is twisted down. Since each angle in G corresponds to a face in $H(G)$, each color of c on the lattice of spanning trees of G must correspond to a color of c' on the isomorphic lattice of d-factors on $H(G)$. Since c' must be an L-coloring, coloring c must also be an L-coloring. □

Lemma 7.9 again allows us to apply our results from Section 4.

Theorem 7.10. *Given a planar graph G with special vertex v^* and special face f^*, consider the distributive lattice of spanning trees on G with special vertex v^* and face f^*. Then each color in S-coloring of the lattice consists of all edges corresponding to moves of the form "the j^{th} swing-down operation at angle eve'" for a fixed angle eve' and positive integer j.*

Proof. By Lemma 7.9, there exists a valid L-coloring in which each color corresponds to a single angle being swung down. By Lemma 4.15, each color in this L-coloring is a union of colors of the S-coloring. Thus, for any angle eve', all colors in the S-coloring corresponding to swing-down operations through eve' must occur in the same order in any complete downward path in the lattice of spanning trees. Thus, each color in the S-coloring corresponding to angle eve' must consist of all moves of the form "the j^{th} move at eve'" for some fixed j. □

By Theorem 7.10, all of the coloring results on S-colorings from Section 4 apply to the lattice of spanning trees. In particular, for any two spanning trees T_1 and T_2 such that $T_1 \geq T_2$, all sequences of swing-down operations must involve swinging down through the same multiset of angles by Theorem 4.11. For any two spanning trees T_1 and T_2, any shortest path from T_1 to T_2 involving swing-up and swing-down operations must involve swinging up and down through the same multiset of angles by Theorem 4.21.

Because of Theorem 7.10, all of the coloring results from Section 4 apply to the poset of spanning trees related by swinging down operations. In particular, we have the following two results:

Corollary 7.11. *Consider a planar graph G with spanning trees T_1 and T_2 such that $T_1 > T_2$ in the poset of spanning trees. Then all sequences of swing-down operations from T_1 to T_2 involve swinging down through the same multiset of angles.*

Proof. The poset of spanning trees is an LLD lattice with an S-coloring as defined in Theorem 7.10. As a consequence of Theorem 4.6, for any S-coloring c, all downward paths from T_1 to T_2 must contain edges of the same colors, with at most one edge of each color. Since each color corresponds to a swinging down move through a particular angle, all sequences of swing-down moves from T_1 to T_2 must involve swinging down through the same multiset of angles. □

Corollary 7.12. *Consider a planar graph G with spanning trees T_1 and T_2. Then any two shortest paths from T_1 to T_2 involving swinging up and swinging down operations must include the same number of swinging up operations through each angle eve', and the same number of swinging down operations through each angle eve'.*

Proof. As the poset of spanning trees is an LLD lattice with an S-coloring as defined in Theorem 7.10, this follows from Lemma 4.19. □

8. S-Colorings and S_n EL-Labelings

Now, we compare LLD lattices and their corresponding S-colorings to another class of lattices that are similarly determined by an edge labeling: supersolvable lattices and their S_n EL-Labelings.

A finite lattice L is supersolvable if it contains a maximal chain M, which together with any other chain in L generates a distributive sublattice [18].

Given a finite graded lattice L of rank n, define an edge labeling λ to be a function from the edges of the Hasse diagram of L to $[n]$. Given a labeling λ, consider a maximal chain $m = (s = s_0 < s_1 < \cdots < s_k = t)$

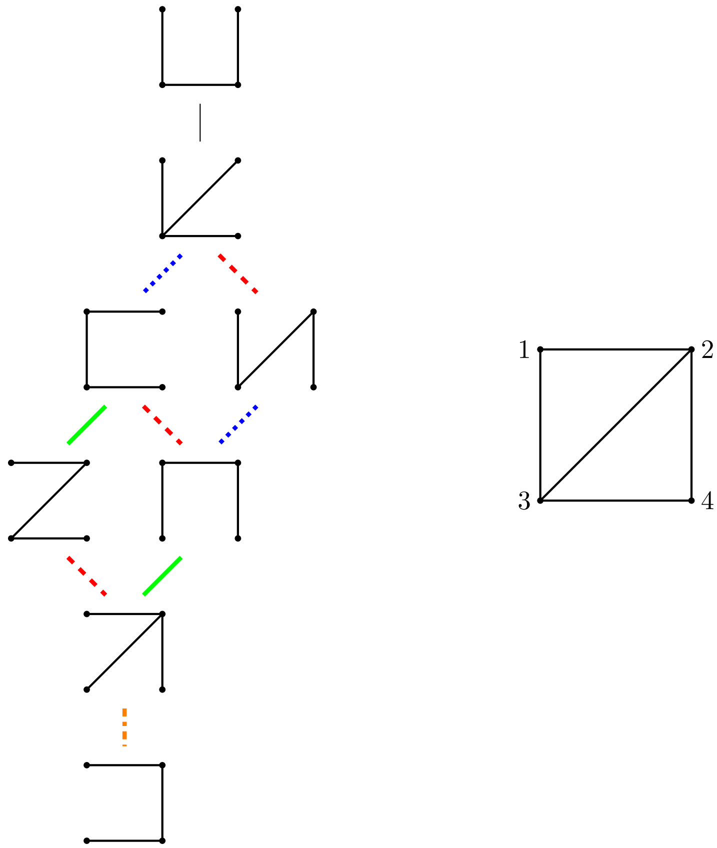

Figure 13. The partial ordering of spanning trees on the diamond graph (K_4 with an edge removed), as shown on the right. Each edge color in the Hasse diagram corresponds to swinging down through a particular angle (the induced process never requires swinging down multiple times through the same angle). Black (thin) corresponds to the angle 423, blue (dotted) corresponds to 321, red (dashed) corresponds to 342, green (thick) corresponds to 132, and orange (dash-dotted) corresponds to 234.

on an interval $[s, t]$. m is **increasing** under labeling λ if $\lambda\big((s_0, s_1)\big) \leq \lambda\big((s_1, s_2)\big) \leq \cdots \leq \lambda\big((s_{k-1}, s_k)\big)$. Let $\leq_L$ denote the lexicographic ordering on maximal chains on $[s, t]$. Given two maximal chains $m = (s = s_0 < s_1 < \cdots < s_k = t)$ and $m' = (s = s'_0 < s'_1 < \cdots < s'_k = t)$, we have $m <_L m'$ if there exists some i such that $\lambda\big((s_i, s_{i+1})\big) < \lambda\big((s'_i, s'_{i+1})\big)$ and for all $j < i$, $\lambda\big((s_j, s_{j+1})\big) = \lambda\big((s'_j, s'_{j+1})\big)$. We provide the following definition:

Definition 8.1. Let P be a finite graded poset of rank n with a $\mathbf{0}$ and a textbf1. An edge labeling $\lambda : E(P) \to [n]$ is an S_n EL-labeling if

(1) Every interval $[s, t]$ has exactly one increasing maximal chain m.
(2) Any other maximal chain m' of $[s, t]$ satisfies $\lambda(m') >_L \lambda(m)$.
(3) The labelings of any maximal chain on $[\mathbf{0}, \mathbf{1}]$ form a permutation of $[n]$.

McNamara shows that that a finite graded lattice of rank n is supersolvable if and only if it admits an S_n EL-labeling [15]. Thus, it is natural to compare the lattices admitting S_n EL-labelings to those admitting S-colorings, as both assign labelings to edges in a similar manner.

We will show that neither class of lattices contains the other. We can construct counterexamples to each inclusion relation.

Example 8.2. There exists a lattice with an S_n EL-labeling that does not admit an S-coloring. A Hasse diagram for such a lattice is shown in Figure 14, along with an S_n EL-labeling.

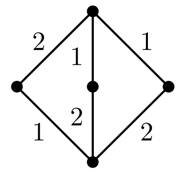

FIGURE 14. A poset P with an S_n EL-labeling, but no S-coloring

Every upward path from $\mathbf{0}$ to $\mathbf{1}$ in poset P contains exactly one 1 and one 2, including one such path that is increasing. All smaller intervals contain at most one edge, and so must satisfy the properties of S_n EL-labelings. Thus, the labeling in Figure 14 is an S_n EL-labeling.

However, poset P does not admit an S-coloring. Any coloring of the Hasse diagram of P that satisfies property (2) of S-coloring must have exactly one color (see Figure 15). However, this would result in every path from $\mathbf{1}$ to $\mathbf{0}$ using the same color twice, thus violating property (1) of S-coloring. Thus, P has no S-coloring, and the set of LLD lattices does not contain the set of supersolvable lattices.

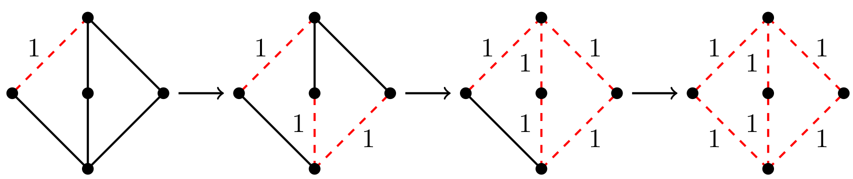

FIGURE 15. Suppose there exists an S-coloring, in which the color of the top-left edge is 1 (represented by red/dashed). Then by property (2) of S-coloring, the color of the bottom-middle and bottom right edges must be 1. This implies that the color of the top-middle and top-right edges is 1, which implies that the color of the bottom-left edge is 1. Thus, the color of all edges must be the same, but this contradicts property (1) of S-coloring.

Example 8.3. There exists an S-colorable lattice that does not admit an S_n EL-labeling. Consider the chip-firing process on the following graph and initial configuration.

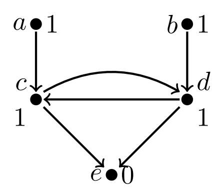

Here, each vertex is labeled with a letter, as well as a number representing the number of chips at that site in the initial configuration. This chip-firing process produces the following configuration poset P. Each edge is labeled with the corresponding vertex being fired:

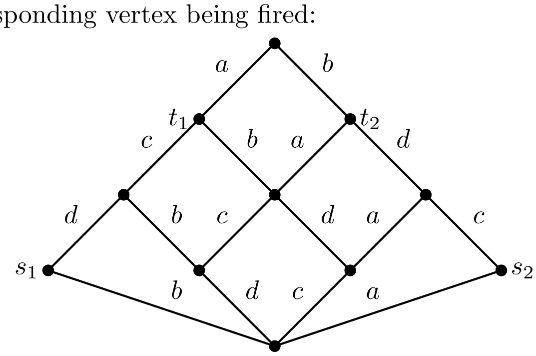

P is S-colorable because it is the poset of configurations for a chip-firing process. Since each site fires at most once, each color in the S-coloring of P corresponds to a specific vertex being fired. Suppose that an S_n EL-labeling exists for poset P. As every interval of length 2 is either a single path or a diamond, any S_n

EL-labeling of P must satisfy the diamond property. Since there are 4 labels and the poset is of rank 4, any S_n EL-labeling must have exactly one label corresponding to each of the sites a, b, c, and d.

The interval $[s_1, t_1]$ consists of a single path with one edge corresponding to site d, and another edge corresponding to site c. Any S_n EL-labeling must have one increasing path from s_1 to t_1, so the label corresponding to d must be less than the label corresponding to c. However, the interval $[s_2, t_2]$ has the same edge labels in the opposite order, so any S_n EL-labeling must have a smaller label for c than for d. This is a contradiction, so no S_n EL-labeling exists for poset P.

As our goal is to produce a Venn diagram relating LLD, supersolvable, and distributive lattices, we include an additional example for completeness. It is known that all distributive lattices are both LLD and supersolvable [15], but we will show that this is a strict inclusion: namely, that there is an LLD, supersolvable lattice that is not distributive.

Example 8.4. Consider the chip-firing process with the following graph and initial configuration.

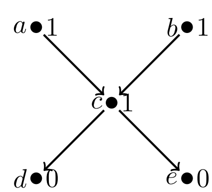

This chip-firing process produces the following poset P of configurations, in which each edge is labeled with the corresponding vertex being fired:

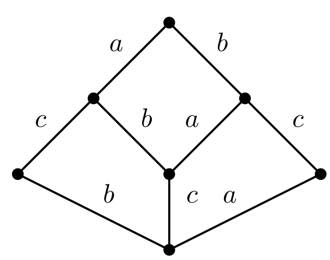

As P is the configuration poset of a finite chip-firing process, it is LLD. Furthermore, the labeling that assigns label 1 to all of the edges corresponding to firing moves at c, label 2 to firing moves at b, and label 3 to firing moves at a, is an S_n EL-labeling. However, P is not a distributive lattice (its dual does not satisfy property (2) of S-colorings, so P is not ULD).

We thus have the following relationship between our relevant classes of lattices: LLD and supersolvable lattices overlap, but neither class includes the other. Furthermore, distributive lattices are included in this overlap, and this is a strict inclusion. These relationships are summarized in the following Venn diagram:

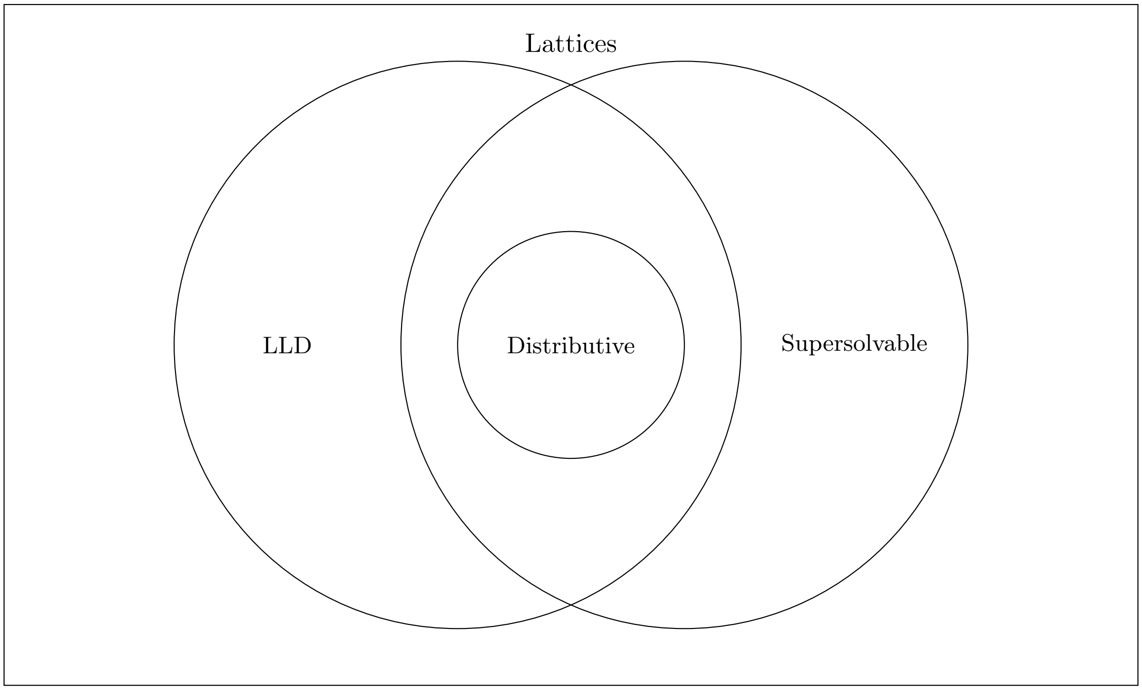

References

[1] R. Anderson, L. Lovász, P. Shor, J. Spencer, E. Tardos, and S. Winograd. Disks, balls, and walls: Analysis of a combinatorial game. *The American Mathematical Monthly*, 96:481–493, 1989.
[2] P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality: An explanation of the 1/f noise. *Phys. Rev. Lett.*, 59:381–384, 1987.
[3] A. Björner and F. Brenti. *Combinatorics of Coxeter groups.* Springer, New York, 2005.
[4] A. Björner, L. Lovász, and P. Shor. Chip-firing games on graphs. *European Journal of Combinatorics*, 12:283–291, 1991.
[5] N. Bourbaki. *Lie groups and Lie algebras. Chapters 4-6.* Springer-Verlag, Berlin, 2002. Translated from the 1968 French original by Andrew Pressley.
[6] Robert Burton and Robin Pemantle. Local characteristics, entropy and limit theorems for spanning trees and domino tilings via transfer-impedances. *The Annals of Probability*, 21:1329–1371, 1993.
[7] S. Felsner and K. Knauer. Uld-lattices and δ-bonds. *Combin. Probab. Comput.*, 18:707–724, 2009.
[8] P. Galashin, S. Hopkins, T. McConville, and A. Postnikov. Root system chip-firing II: Central-firing. *International Mathematics Research Notices*, 2017.
[9] S. Hopkins, T. McConville, and J. Propp. Sorting via chip-firing. *Electronic Journal of Combinatorics*, 24, 2016.
[10] J. Humphreys. *Introduction to Lie algebras and representation theory.* Springer-Verlag, New York-Berlin, 1972.
[11] C. Klivans and P. Liscio. Confluence in labeled chip-firing. *Journal of Combinatorial Theory, Series A*, 186, 2022.
[12] P. Liscio. *Confluence in Chip-Firing and Related Combinatorial Processes.* PhD thesis, Brown University, 2022.
[13] László Lovász. *Combinatorial problems and exercises*, volume 361. American Mathematical Soc., North Holland, 1993.
[14] C. Magnien, H. D. Phan, and L. Vuillon. Characterization of lattices induced by (extended) chip firing games. *Discrete Mathematics and Theoretical Computer Science Proceedings*, pages 229–244, 2001.
[15] P. McNamara. El-labelings, supersolvability and 0-hecke algebra actions on posets. *J. Combin. Theory Ser. A*, 101:69–89, 2003.
[16] J. Propp. Lattice structure for orientations of graphs, 1993. arXiv:math/0209005.
[17] N. C. Saldanha, C. Tomei, M. A. Casarin Jr., and D. Romualdo. Spaces of domino tilings. *Discrete Comput. Geom.*, 14:207–233, 1995.
[18] R. Stanley. Supersolvable lattices. *Algebra Universalis*, 2:197–217, 1972.
[19] H. N. V. Temperley. In combinatorics (london math. soc. lecture note series# 13), 1974.
[20] W. Thurston. Conway's tiling groups. *American Mathematical Monthly*.