

ULD-Lattices and Δ -Bonds

STEFAN FELSNER[†] and KOLJA B. KNAUER[‡]

Institut für Mathematik, Technische Universität Berlin, Germany
(e-mail: {felsner, knauer}@math.tu-berlin.de)

Received 4 July 2008; revised 13 February 2009; first published online 1 June 2009

We provide a characterization of upper locally distributive lattices (ULD-lattices) in terms of edge colourings of their cover graphs. In many instances where a set of combinatorial objects carries the order structure of a lattice, this characterization yields a slick proof of distributivity or UL-distributivity. This is exemplified by proving a distributive lattice structure on Δ -bonds with invariant circular flow-difference. This instance generalizes several previously studied lattice structures, in particular, c -orientations (Propp), α -orientations of planar graphs (Felsner, resp. de Mendez) and planar flows (Khuller, Naor and Klein). The characterization also applies to other instances, *e.g.*, to chip-firing games.

1. Introduction

The concept of upper locally distributive lattices (ULDs) – and its dual, lower locally distributive lattices (LLDs) – has appeared under several different names, *e.g.*, locally distributive lattices (Dilworth [10]), meet-distributive lattices (Jamison [15, 16], Edelman [11], Björner and Ziegler [7]), and locally free lattices (Nakamura [21]). Following Avann [2], Monjardet [20], Stern [26] and others, we call them ULDs. The reason for the frequent reappearance of the concept is that there are many instances of ULDs, *i.e.*, sets of combinatorial objects that can be *naturally* ordered to form a ULD.

ULDs were first investigated by Dilworth [9], and many different lattice-theoretical characterizations of ULDs are known. For a survey on the work up to 1990 we refer to Monjardet [20]. We use Dilworth's original definition, as follows.

Definition 1. Let (P, \leq) be a poset. P is an *upper locally distributive lattice* (ULD) if P is a lattice and each element has a unique minimal representation as a meet of meet-irreducibles, *i.e.*,

[†] Partially supported by Deutsche Forschungs-Gemeinschaft (DFG) Fe 340/7-1.

[‡] Supported by the Research Training Group *Methods for Discrete Structures* (DFG-GRK 1408).

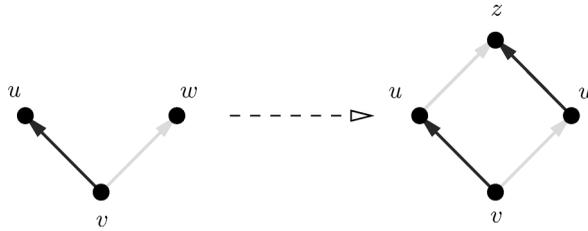


Figure 1. The completion property of U-colourings.

there is a mapping $M : P \rightarrow \mathcal{P}(\{m \in P : m \text{ is meet-irreducible}\})$ with the properties:

- $x = \bigwedge M_x$ (representation),
- $x = \bigwedge A$ implies $M_x \subseteq A$ (minimal).

Let $D = (V, A)$ be a directed graph. An arc colouring c of D is an *U-colouring* if, for every $u, v, w \in V$ with $u \neq w$ and $(v, u), (v, w) \in A$, we have:

(U₁) $c(v, u) \neq c(v, w)$,

(U₂) there is a $z \in V$ and arcs $(u, z), (w, z)$ such that $c(v, u) = c(w, z)$ and $c(v, w) = c(u, z)$ (see Figure 1).

Definition 2. A finite poset (P, \leq) is called a *U-poset* if the arcs of the cover graph D_P of P admit a U-colouring.

The characterization of ULDs in Section 2 has two parts.

Theorem 1.1.

- (a) *If D is a finite, acyclic digraph admitting a U-colouring, then D is a cover graph. Hence, the transitive closure of D is a U-poset.*
- (b) *Upper locally distributive lattices are exactly the U-posets with a global minimum.*

The duals, in the sense of order reversal, to U-colouring, U-poset and ULD are L-colouring, L-poset and LLD, respectively. The characterization of LLDs dual to Theorem 1.1 allows easy proofs that the inclusion orders on the following combinatorial structures are lower locally distributive lattices:

- subtrees of a tree (Boulaye [8]),
- convex subsets of posets (Birkhoff and Bennett [3]),
- convex subgraphs of acyclic digraphs; here a set C is convex if $x, y \in C$ implies that all directed (x, y) -paths are in C (Pfaltz [22]).

These combinatorial structures can also be seen as *convex sets of an abstract convex geometry*. This is no coincidence, as every LLD is isomorphic to the inclusion order on the convex sets of an abstract convex geometry (Edelman [11]). The L-colourings of the cover-relations of LLDs are EL-shellings, in the language of Björner [4]. They give rise to what is called compatible orderings of the ground set (see [12]). In a more recent work, Ardila and Maneva [1] provide yet another characterization of LLDs.

Dual to the representations of LLDs as convex sets, every ULD is isomorphic to the inclusion order on the independent sets of an antimatroid (see [19]). In the present paper we represent ULDs as inclusion orders on *multisets* (of colours). In a sense this corresponds to the concepts of *antimatroids with repetition* [7] and *locally free, permutable, left-hereditary languages* [6]. There are in fact objects like chip-firing games that naturally lead to ULD representations by multisets and not by sets: see Section 4.5.

In Section 3 we deal with distributive lattices arising from orientations of graphs. To prove distributivity we use the following well-known characterization: *Distributive lattices are exactly those lattices that are both ULD and LLD* (Theorem 2.13).

Let $D = (V, A)$ be a connected directed graph with upper and lower integral edge capacities $c_u, c_l : A \rightarrow \mathbb{Z}$. We are interested in maps $x : A \rightarrow \mathbb{Z}$ such that $c_l(a) \leq x(a) \leq c_u(a)$ for all $a \in A$. The *circular flow-difference* of x on a cycle C with a prescribed direction is

$$\delta(C, x) := \sum_{a \in C^+} x(a) - \sum_{a \in C^-} x(a).$$

Note that the circular flow-differences $\delta(C, x)$ on the cycles of a basis of the cycle space uniquely determines the flow-difference of x on all cycles of the graph.

For a given $\Delta \in \mathbb{Z}^{\mathcal{C}}$ we consider the set

$$\mathcal{B}_\Delta(D, c_l, c_u) := \{c_l \leq x \leq c_u \mid \delta(C, x) = \Delta_C \text{ for all } C \in \mathcal{C}\},$$

and call it the set of Δ -bonds on (D, c_l, c_u) , the name *bonds* arising because $\mathcal{B}_0(D, c_l, c_u)$ is orthogonal to the space of flows of D . We introduce an order on Δ -bonds with prescribed circular flow-difference, *i.e.*, on the elements of $\mathcal{B}_\Delta(D, c_l, c_u)$.

Theorem 1.2. $\mathcal{B}_\Delta(D, c_l, c_u)$ carries the structure of a distributive lattice.

Theorem 1.2 is restated in Section 3 in a more precise version as Theorem 3.3. The power of this result is exemplified in Section 4, where we show that several previously studied distributive lattices can be recognized as special cases of Δ -bonds. This is shown in the following cases:

- lattices of c -orientations of graphs (Propp [23]),
- lattices of flow in planar graphs (Khuller, Naor and Klein [17]),
- lattices of α -orientations of planar graphs (Felsner [13]).

In Section 4.5 we discuss the *chip-firing game* on directed graphs. Important properties of this game can be proved in the context of U-posets and ULD lattices.

2. A characterization of upper locally distributive lattices

In this section we prove Theorem 1.1. This new characterization of ULD lattices can clearly be used to show equivalence to any of the many known characterizations of ULDs: see [20]. In our proof we will establish the equivalence to the original definition given by Dilworth [9]. Finally, we add a proof of the known fact that a lattice which is both ULD and LLD is actually distributive.

The following lemma is the main tool for the proof.

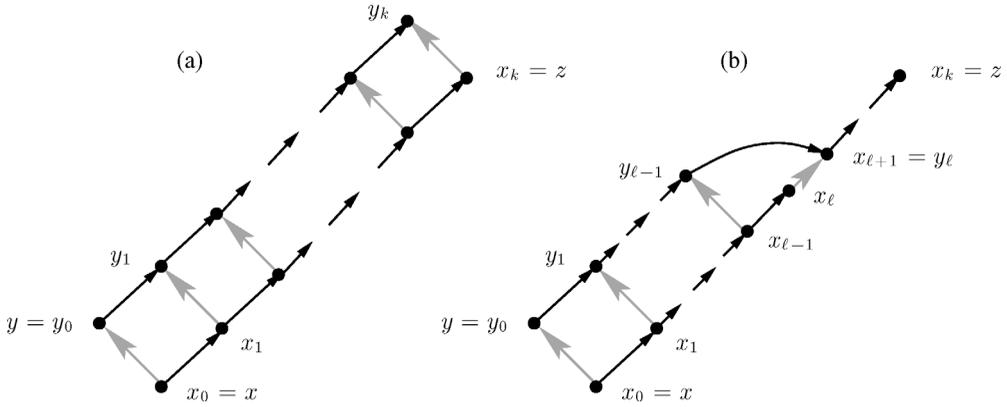


Figure 2. Illustration of Lemma 2.1.

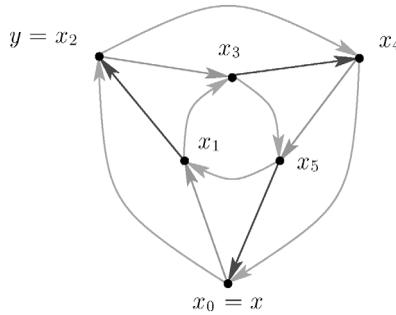


Figure 3. A digraph with a U-colouring. Choosing $p = x_0, \dots, x_5$ and $y = x_2$, we get $y_i = x_{i+2 \pmod 6}$ for $i = 0, \dots, 6$.

Lemma 2.1. *Let $D = (V, A)$ be a digraph with a U-colouring c . If (x, y) is an arc and $p = x_0, \dots, x_k$ is a directed path from $x = x_0$ to $z = x_k$, then there is a sequence y_0, \dots, y_ℓ such that $(x_i, y_i) \in A$ and $c(x_i, y_i) = c(x, y)$ for $i = 1, \dots, \ell$ and either $\ell = k$ or $\ell < k$ and $y_\ell = x_{\ell+1}$. Figure 2 illustrates the two cases (a) and (b). Case (b) occurs if and only if there is an edge $(x_\ell, x_{\ell+1})$ on p with $c(x, y) = c(x_\ell, x_{\ell+1})$.*

Proof. We recursively apply rule U_2 to edges (x_i, y_i) and (x_i, x_{i+1}) to define a vertex y_{i+1} with edges (y_i, y_{i+1}) and (x_{i+1}, y_{i+1}) such that $c(x_i, y_i) = c(x_{i+1}, y_{i+1})$. The iteration ends either if $i = k$ (case (a)), or if the two edges needed for the next application of the rule are the same, i.e., $y_i = x_{i+1}$ (case (b)). In this case $c(x_i, x_{i+1}) = c(x_{i-1}, y_{i-1}) = c(x, y)$, i.e., there is an edge on the path p whose colour equals the colour of edge (x, y) . Rule U_1 implies that case (b) occurs whenever there is an edge on p whose colour equals the colour of edge (x, y) . \square

Remark 1. The proof does not imply that $y_i \neq x_j$ in all cases, as is suggested by Figure 2. An example is given in Figure 3. From the analysis below it follows that in all *bad* cases D is infinite or not acyclic.

From now on we assume that $D = (V, A)$ is a finite, connected and acyclic digraph with a U-colouring c . The assumptions imply that the transitive closure of D is a finite poset P_D . From the next two propositions it will follow that D is transitively reduced, *i.e.*, the cover graph of P_D . Hence, P_D is a U-poset.

Proposition 2.2. *There is a unique sink in D .*

Proof. Since D is acyclic and finite it has a sink. Suppose that there are two sinks s_0 and s_1 . Let p be a shortest (s_0, s_1) -path in the underlying undirected graph. In D the first and the last edge of p are oriented towards the sources. Hence, there is a last vertex x on p which is a source on p , *i.e.*, the final part of p is $y \leftarrow x \rightarrow x_1 \rightarrow \dots \rightarrow s_1$. With $z = s_1$ this is the precondition of Lemma 2.1, but both outcomes yield a contradiction: (a) is impossible because s_1 is a sink; (b) is impossible because it implies that there is a shorter (s_0, s_1) -path than p . \square

We define the *colourset* $c(p)$ of a directed path p as the multiset of colours used on edges of p .

Proposition 2.3. *If D has a unique sink s and p, p' are directed (x, z) -paths, then $c(p) = c(p')$. In particular this allows us to define the colourset $c(x)$ of a vertex x as the colourset of any (x, s) -path.*

Proof. First assume that $z = s$. In a top-down induction we show that the coloursets of all (x, s) -paths are equal.

Assume that we know that $c(v)$ is well defined for all vertices v that are accessible from x via directed paths in D . Let p and p' be (x, s) -paths with different initial edges. Let y be the successor of x on p' and let a be the colour of (x, y) . The edge (x, y) and the path p form the precondition for Lemma 2.1. Since $z = s$ is the sink, the situation has to be as in (b). Let p'' be the path $y \rightarrow y_1 \rightarrow \dots \rightarrow y_\ell \rightarrow \dots \rightarrow s$. The assumption for y yields $c(p') = a \oplus c(y) = a \oplus c(p'')$. The colouring rule U_2 implies that $c(x_i, x_{i+1}) = c(y_i, y_{i+1})$ for all $i = 0, \dots, \ell - 1$ and $a = c(x, y) = c(x_\ell, x_{\ell+1})$; therefore, $a \oplus c(p'') = c(p)$. This shows that $c(x) = c(p') = c(p)$ is well defined.

If p is a directed (x, z) -path, then the concatenation $p \circ q$ with a (z, s) -path q yields an (x, s) -path; hence $c(x) = c(p \circ q) = c(p) \oplus c(q) = c(p) \oplus c(z)$. This shows that $c(p) = c(x) \ominus c(z)$ depends only on the end-vertices. \square

Since the colourset of a directed (x, z) -path in D depends only on the end-vertices, we also know that all (x, z) -paths have the same length. This implies that D is transitively reduced. We have thus shown the following, which is slightly stronger than statement (a) of Theorem 1.1.

Corollary 2.4. *If $D = (V, A)$ is a finite, connected and acyclic digraph with a U-colouring c , then D is a cover graph and its transitive closure is a U-poset P_D with a rank function and a $\mathbf{1}$.*

Let P be a U-poset with a global minimum $\mathbf{0}$. Define a mapping $\gamma : P \rightarrow \mathbb{N}^k$, where k is the number of colours of the U-colouring. Assuming that the colours used by c are $1, \dots, k$, the

i th component of $\gamma(x)$ is the multiplicity of colour i on any $(\mathbf{0}, x)$ -path in the cover graph D_P . Identifying vectors in \mathbb{N}^k with multisets of colours, we have $\gamma(x) = c(\mathbf{0}) - c(x)$.

Lemma 2.5. *Let P be a U -poset with a global minimum $\mathbf{0}$. The mapping $\gamma : P \rightarrow \mathbb{N}^k$ is an order-preserving embedding of P into the dominance order on \mathbb{N}^k .*

Proof. The implication from $y \leq_P z$ to $\gamma(y) \leq \gamma(z)$ follows from the fact that extending a path requires more colours. In particular, the number of edges of colour i on a $(\mathbf{0}, z)$ -path is at least as large as on a $(\mathbf{0}, y)$ -path, i.e., $\gamma_i(y) \leq \gamma_i(z)$. In fact $y <_P z$ implies $\gamma(y) \neq \gamma(z)$.

For the converse suppose $\gamma(y) \leq \gamma(z)$ but $y \not\leq_P z$. From the first part and $y \neq z$ we know $z \not\leq_P y$. Let x be maximal with the property $x \leq_P z$ and $x \leq_P y$. Consider the first edge (x, y') on an (x, y) -path in D_P , and let p be an (x, z) -path. This is a situation for Lemma 2.1. Since the colour of edge (x, y') also occurs on p , we are in case (b). This case, however, is impossible because $y' \leq_P z$, $y' \leq_P y$ and $y' \geq x$ contradict the choice of x . □

It is easy to see that an embedding $\gamma : P \rightarrow \mathbb{N}^k$ yields an embedding $\gamma' : P \rightarrow \{0, 1\}^{k'}$. Since generally $k' > k$, embedding into the grid instead of the cube might be useful.

Lemma 2.6. *Let P and γ be as above. For all $z, y \in P$ there is a $w \in P$ with $\gamma(w) = \gamma(y) \vee \gamma(z)$ where \vee is the componentwise maximum.*

Proof. For any fixed y we proceed with top-down induction. Given z , consider a maximal x with the property $x \leq_P z$ and $x \leq_P y$. Let (x, y') be the first edge on an (x, y) -path in D_P and let p be an (x, z) -path. Case (b) of Lemma 2.1 is impossible because y' would have prevented us from choosing x . Hence, we are in case (a) and there is a z' covering z such that the edges (z, z') and (x, y') have the same colour i and, moreover, the path from x to z has no edge of colour i . Induction implies that there is a w' such that $\gamma(w') = \gamma(z') \vee \gamma(y)$. Since $\gamma(z') = \gamma(z) + e_i$ and for the i th component we have $\gamma_i(z) = \gamma_i(x) < \gamma_i(y)$, we can conclude $\gamma(z') \vee \gamma(y) = \gamma(z) \vee \gamma(y)$, i.e., w' may also serve as w . □

Proposition 2.7. *If a U -poset has a global minimum $\mathbf{0}$, then it is a lattice.*

Proof. The mapping γ is an order embedding of P into \mathbb{N}^k (Lemma 2.5) and the image $\gamma(P)$ is join-closed (Lemma 2.6). Together this implies that P has unique least upper covers (joins exist). Since P has a $\mathbf{0}$ there is a lower cover for every pair of elements. Uniqueness for greatest lower covers (meets) follows from the existence of unique joins. □

The next goal is to show that every element of P has a unique minimal representation as a meet of meet-irreducibles. Let $C(x)$ be the set of colours of the edges emanating from x in the cover graph. In the next lemma we associate a meet-irreducible element with every colour $i \in C(x)$.

Lemma 2.8. *Let P be a U -poset with a $\mathbf{0}$ and let $x \neq \mathbf{1}$ be an element of P . For every $i \in C(x)$ there is a unique maximal element y_i such that:*

- $y_i \geq x$ and

- $\gamma_i(y_i) = \gamma_i(x)$.

The element y_i is meet-irreducible and $\gamma_j(y_i) > \gamma_j(x)$ for all $j \in C(x) \setminus \{i\}$.

Proof. Let $i \in C(x)$ and consider the set $S_i(x)$ of all $y \geq x$ with $\gamma_i(y) = \gamma_i(x)$. The set $S_i(x)$ contains x ; hence it is non-empty, and by Lemma 2.6 it contains a unique maximal element y_i . The element y_i is meet-irreducible, otherwise we could find a successor of y_i in $S_i(x)$. For every $j \in C(x)$ there is an element x_j with $\gamma(x_j) = \gamma(x) + e_j$. For $j \neq i$ the element x_j is in $S_i(x)$; hence $\gamma(x_j) \leq \gamma(y_i)$ and $\gamma_j(x) < \gamma_j(x_j) \leq \gamma_j(y_i)$. □

Proposition 2.9. A U-poset with a global minimum $\mathbf{0}$ is an upper locally distributive lattice.

Proof. We claim that $M_x = \{y_i : i \in C(x)\}$ is the unique minimal set of meet-irreducibles with $x = \bigwedge M_x$.

Let z be any lower bound for M_x , i.e., an element with $z \leq y_i$ for all $i \in C(x)$. Since γ is order-preserving and \mathbb{N}^k is closed under taking meets, we have $\gamma(z) \leq \bigwedge \{\gamma(y_i) : i \in C(x)\}$. From $\gamma_i(y_i) = \gamma_i(x)$ it follows that $\bigwedge \{\gamma(y_i) : i \in C(x)\} = \gamma(x)$. Since γ is order-preserving, this implies $z \leq x$, i.e., x is the unique maximal lower bound for M_x and the notation $x = \bigwedge M_x$ is justified.

It remains to show that the representation $x = \bigwedge M_x$ is the unique minimal representation of x as a meet of meet-irreducibles. Let $i \in C(x)$ and consider a set M of meet-irreducibles with $y_i \notin M$. It is enough to show that $x \neq \bigwedge M$. If M contains a y with $x \not\leq y$, then $x \neq \bigwedge M$ is obvious. Consider the set $S_i(x)$ from the proof of Lemma 2.8; every element $y \neq y_i$ in this set is contained in an (x, y_i) -path p that contains no i -coloured edge. Lemma 2.1 implies that there is an i -coloured edge leaving y and, together with the edge leaving y on p , this implies that y is not meet-irreducible. Hence $M \cap S_i(x) = \emptyset$. All $y > x$ with $y \notin S_i(x)$ satisfy $y \geq x_i$. Given our assumptions on M , this implies that x_i with $\gamma(x_i) = \gamma(x) + e_i$ is a lower bound on M , i.e., $x \neq \bigwedge M$. □

From what we have shown so far we obtain the following criterion for ULD lattices.

Theorem 2.10. If D is a finite, acyclic digraph with a unique source $\mathbf{0}$ and there is a U-colouring of the arcs of D , then the transitive closure of D is an upper locally distributive lattice.

To complete the proof of Theorem 1.1 it remains to show that every ULD has a representation as a U-poset, i.e., we have to present a U-colouring of its cover graph.

Theorem 2.11. The cover graph of every finite ULD lattice admits a U-colouring.

Consider the mapping \overline{M} that takes an element x of P to the set \overline{M}_x of all meet-irreducible elements that are at least as large as x . The definition of meet-irreducible implies that $x = \bigwedge \overline{M}_x$ for all x , i.e., the set \overline{M}_x uniquely characterizes x . Moreover, $x \leq y$ if and only if $\overline{M}_x \supseteq \overline{M}_y$.

On the basis of the mappings M and \overline{M} we will define a U-colouring of the cover relations of P . As colours we use the meet-irreducible elements of P .

Lemma 2.12. *Let P be a ULD lattice. A comparability $x < y$ is a cover if and only if $|\overline{M}_x \setminus \overline{M}_y| = 1$.*

Proof. An element z with $x < z < y$ satisfies $\overline{M}_y \subsetneq \overline{M}_z \subsetneq \overline{M}_x$, hence $|\overline{M}_x \setminus \overline{M}_y| \geq 2$.

Let $x < y$ and suppose that $|\overline{M}_x \setminus \overline{M}_y| \geq 2$. Since $\bigwedge M_x < \bigwedge \overline{M}_y$, there has to be some $m \in M_x \setminus \overline{M}_y$. Let $z = \bigwedge(\overline{M}_x - m)$; we claim that $x < y < z$. From $\overline{M}_x - m \subset \overline{M}_x$ it follows that $z \geq x$, and since $m \in M_x$ and $m \notin M_z$, Definition 1 implies that $z < x$. Since $(\overline{M}_x - m) \supseteq \overline{M}_y$, we have $z \leq y$. Let m' be an element with $m \neq m' \in \overline{M}_x \setminus \overline{M}_y$. It follows that $m' \in \overline{M}_z$ and $m' \notin \overline{M}_y$. Therefore $z \neq y$, and we have shown that $x < z < y$, i.e., the pair x, y is not in a cover relation. □

To a cover relation $x < y$ we assign the unique meet-irreducible in $\overline{M}_x \setminus \overline{M}_y$ as its colour. Note that this meet-irreducible is a member of M_x . To verify that this is a U-colouring we have to check the two properties U_1 and U_2 .

Claim 1. *The colouring satisfies U_1 .*

Proof. Let $x < y_1$ and $x < y_2$ be two cover relations. Since x is the meet of y_1 and y_2 , we have the representation $x = \bigwedge(\overline{M}_{y_1} \cup \overline{M}_{y_2})$ of x as a meet of irreducibles; hence $M_x \subseteq \overline{M}_{y_1} \cup \overline{M}_{y_2}$. If both covers had the same colour m , then $m \in M_x$, but $m \notin \overline{M}_{y_1} \cup \overline{M}_{y_2}$, a contradiction. □

Claim 2. *The colouring satisfies U_2 .*

Proof. Let $x < y_1$ and $x < y_2$ be two cover relations such that $x < y_1$ has colour m_1 and $x < y_2$ has colour m_2 , i.e., $\overline{M}_{y_i} = \overline{M}_x - m_i$. Consider $z = \bigwedge(\overline{M}_x - m_1 - m_2)$. Since z is representable as a meet of elements from \overline{M}_{y_i} we know that $z \geq y_i$ for $i = 1, 2$. Since y_1 and y_2 both cover x it follows that $z \neq y_1, y_2$, hence $z > y_i$ for $i = 1, 2$. From $\overline{M}_x - m_1 - m_2 \subseteq \overline{M}_z \subset \overline{M}_{y_i} = \overline{M}_x - m_i$ it follows that $|\overline{M}_{y_i} \setminus \overline{M}_z| = 1$. Lemma 2.12 implies that z covers each y_i , and the labels of these covers are as required. □

In many applications of the characterization of ULDs the lattice in question is actually distributive. Such a situation is the topic of the next section. To make the paper self-contained we prove the following folklore result.

Theorem 2.13. *If a finite, acyclic and connected digraph D admits a U- and an L-colouring, then D is the cover graph of a distributive lattice P_D . Moreover, the colourings yield an explicit cover-preserving embedding $P_D \hookrightarrow \mathbb{N}^k$, where k is the number of colours.*

Proof. Corollary 2.4 and its dual imply that D is the cover graph of a poset P_D with $\mathbf{0}$ and $\mathbf{1}$. Hence, by Proposition 2.7, P_D is a lattice.

Let c_U and c_L be a U- and an L-colouring of D . Consider the colouring $c = c_U \times c_L$. The claim is that c is both a U- and an L-colouring of D . The rule U_1 and its dual L_1 are immediately inherited from the corresponding rules for c_U and c_L . Whenever there is a diamond $x < y_1$,

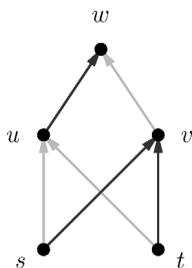


Figure 4.

$x < y_2, y_1 < z, y_2 < z$, the colours of a pair of (x, y_i) and (y_j, z) in c_U and c_L coincide. This implies rules U_2 and L_2 for c .

Consider the order embedding $\gamma : P_D \rightarrow \mathbb{N}^k$ that is based on the colouring c . By Lemma 2.6, γ is compatible with joins; the dual implies that γ is compatible with meets. Therefore γ is a lattice embedding, i.e., P is a sublattice of the distributive lattice \mathbb{N}^k , and hence P is itself distributive. □

Remark 2. Let D be a digraph with a U-colouring. We need acyclicity, connectedness, finiteness and the unique source to conclude that D corresponds to a finite ULD lattice. We feel that among these conditions the *unique source* has a somewhat artificial flavour. Abstaining on this condition it can be shown (along the lines of our proof) that the corresponding poset P is a join-semilattice with the property that, for all $x \in P$, there is a unique minimal set M_x of meet-irreducibles such that x is a maximal lower bound for M_x . Figure 4 shows a small example, in this case $M_s = M_t = \{u, v\}$.

3. The lattice of Δ -bonds

Recall the setting from the Introduction. The data are a directed multigraph $D = (V, A)$ with upper and lower integral edge capacities $c_u, c_l : A \rightarrow \mathbb{Z}$ and a number Δ_C for each cycle C of D . We are interested in the set $\mathcal{B}_\Delta(D, c_l, c_u)$ of Δ -bonds for these data, i.e., the set of maps $x : A \rightarrow \mathbb{Z}$ such that:

- (D₁) $c_l(a) \leq x(a) \leq c_u(a)$ for all $a \in A$ (capacity constraints),
- (D₂) $\Delta_C = \sum_{a \in C^+} x(a) - \sum_{a \in C^-} x(a) = \delta(C, x)$ for all C (circular flow conditions).

Throughout the discussion we shall assume that the data (D, c_l, c_u, Δ) are such that the set of corresponding Δ -bonds is non-empty. Moreover, we want to simplify matters by concentrating on connected graphs and getting rid of *rigid* edges: these are edges $a \in A$ with $x(a) = y(a)$ for all pairs x, y of Δ -bonds.

Let a be a rigid edge of D and let D/a be obtained from D by contracting edge a . Since we allow multiple edges and loops, the cycles in D/a and in D are in bijection. Let C/a be the cycle in D/a corresponding to C in D . Define $\Delta'_{C/a} = \Delta_C$ if $a \notin C$ and $\Delta'_{C/a} = \Delta_C - x(a)$ if $a \in C^+$ and $\Delta'_{C/a} = \Delta_C + x(a)$ if $a \in C^-$. These settings yield the bijection that proves the following.

Lemma 3.1. $\mathcal{B}_\Delta(D, c_l, c_u) \cong \mathcal{B}_{\Delta'}(D/a, c_l, c_u)$.

Given data $(D^1, c_l^1, c_u^1, \Delta^1)$ and $(D^2, c_l^2, c_u^2, \Delta^2)$, there is an obvious extension to a union structure (D, c_l, c_u, Δ) where D is the union of graphs and the c_l, c_u, Δ are concatenations of vectors. Since Δ -bonds factor into a Δ_1 - and a Δ_2 -bond, we have the following result.

Lemma 3.2. $\mathcal{B}_\Delta(D, c_l, c_u) \cong \mathcal{B}_{\Delta^1}(D^1, c_l^1, c_u^1) \times \mathcal{B}_{\Delta^2}(D^2, c_l^2, c_u^2)$.

The data (D, c_l, c_u, Δ) are *reduced* if D is connected and there is no rigid edge. Henceforth we will assume that any given set of data is reduced.

With a partition (U, \bar{U}) of the vertices V of D , we consider the cut $S = S[U] \subset A$. The *forward edges* S^+ of S are those $a \in A$ directed from U to \bar{U} ; *backward edges* S^- of S are directed from \bar{U} to U .

For $x : A \rightarrow \mathbb{Z}$ and a subset $U \subset V$, we define $y = \text{push}_U(x)$ such that $y(a) = x(a) + 1$ for all $a \in S^+[U]$, $y(a) = x(a) - 1$ for all $a \in S^-[U]$, and $y(a) = x(a)$ for all $a \notin S[U]$. We say that $y = \text{push}_U(x)$ is obtained by *pushing* U in x .

Fix an arbitrary vertex v_0 in D as the *forbidden vertex*. For $x, y \in \mathcal{B}_\Delta(D, c_l, c_u)$ define $x \leq y$ if y can be reached from x via a sequence of pushes at sets U_i , such that $v_0 \notin U_i$ for all i and all intermediate states in the sequence are in $\mathcal{B}_\Delta(D, c_l, c_u)$.

Below, in Corollary 3.6, we show that the relation ' \leq ' makes $\mathcal{B}_\Delta(D, c_l, c_u)$ into a partial order $\mathcal{P}_\Delta(D, c_l, c_u)$.

The main result of this section is Theorem 1.2, which can now be stated more precisely.

Theorem 3.3. *The order $\mathcal{P}_\Delta(D, c_l, c_u)$ on Δ -bonds is a distributive lattice.*

It would have been more precise to write $\mathcal{P}_\Delta(D, c_l, c_u, v_0)$ in the theorem because the actual lattice depends on the choice of the forbidden vertex. Different forbidden vertices yield different lattices on the same ground set. The result of Lemma 3.2 carries over to the lattices, i.e., if D^1 and D^2 are connected, then $\mathcal{P}_\Delta(D, c_l, c_u, \{v_1, v_2\}) \cong \mathcal{P}_{\Delta^1}(D^1, c_l^1, c_u^1, v_1) \times \mathcal{P}_{\Delta^2}(D^2, c_l^2, c_u^2, v_2)$.

The next lemma describes the condition for a legal push, i.e., a push that transforms a Δ -bond into a Δ -bond.

Lemma 3.4. *If x is in $\mathcal{B}_\Delta(D, c_l, c_u)$ and $S = S[U]$ is a cut such that $x(a) < c_u(a)$ for all $a \in S^+$ and $x(a) > c_l(a)$ for all $a \in S^-$, then $y = \text{push}_U(x)$ is also in $\mathcal{B}_\Delta(D, c_l, c_u)$.*

Proof. The assumption on $c_u(a)$ and $c_l(a)$ implies that $y = \text{push}_U(x)$ respects the capacity constraints. From the orthogonality of the cycle space and the bond space of D , it follows that $\delta(C, x) = \delta(C, y)$ for all cycles C , i.e., y satisfies the circular flow conditions. □

Lemma 3.5. *If $x \leq y$, then y can be obtained from x by a sequence of vertex pushes, i.e., at cuts $S[v]$ with $v \neq v_0$.*

Proof. It is enough to show that a single push of a set U can be replaced by a sequence of vertex pushes. Let $y = \text{push}_U(x)$. We show that there is a vertex $w \in U$ such that $x' = \text{push}_w(x)$ is a Δ -bond and $y = \text{push}_{U \setminus w}(x')$. This implies the result via induction on the size of U .

Choose any $w \in U$. If pushing w is legal, *i.e.*, $x(a) < c_u(a)$ for all $a \in S[w]^+$ and $x(a) > c_l(a)$ for all $a \in S[w]^-$, then let $x' = \text{push}_w(x)$. Note that pushing $U \setminus w$ in x' is legal. Indeed, for $a \in S[U \setminus w] \cap S[U]$ we have $x'(a) = x(a)$ and all $a \in S[U \setminus w] \setminus S[U]$ are incident to w . If such an a is in $S[U \setminus w]^+$, then it is in $S[w]^-$ and $x'(a) = x(a) - 1 \leq c_u(a) - 1$, and hence $x'(a) < c_u(a)$. The case for $a \in S[U \setminus w]^-$ is symmetric.

It remains to show that there is a $w \in U$ such that pushing w is legal. Choose any $w_1 \in U$. If $S[w_1]$ is not legal, then there is an incident edge $a_1 \in S[w_1]^+$ with $x(a_1) = c_u(a_1)$ or $a_1 \in S[w_1]^-$ with $x(a_1) = c_l(a_1)$. Let w_2 be the second vertex of a_1 and note that $w_2 \in U$. If pushing w_2 is not legal, then there is an incident edge a_2 obstructing the push, and so on. This yields a sequence $w_1, a_1, w_2, a_2, \dots$. Either the sequence ends in a vertex w_j , which is legal for pushing, or it closes into a cycle. Assume that there is a cycle C such that $x(a) = c_u(a)$ for all $a \in C^+$ and $x(a) = c_l(a)$ for all $a \in C^-$. The condition implies $\Delta_C = \delta(C, x) = \sum_{a \in C^+} c_u(a) - \sum_{a \in C^-} c_l(a)$. It follows that every Δ -bond y has $y(a) = x(a)$ for all $a \in C$, *i.e.*, the edges in C are rigid. However, the data (D, c_u, c_l, Δ) are assumed to be reduced, *i.e.*, there are no rigid edges. Hence, there must be a $w \in U$ that is legal for pushing. □

In the case of a general set of data (D, c_u, c_l, Δ) with rigid edges, we would have to allow the pushing of U if and only if U is the vertex set of a connected component of rigid edges.

Corollary 3.6. *The relation \leq on $\mathcal{B}_\Delta(D, c_l, c_u)$ is acyclic, *i.e.*, it is an order relation.*

Proof. Otherwise we could linearly combine vertex cuts $S[v]$, with $v \neq v_0$, to zero. But these vertex cuts are a basis of the bond space. □

Pushes at vertices correspond to cover relations, *i.e.*, edges of the cover graph, of $\mathcal{P}_\Delta(D, c_l, c_u)$ (Lemma 3.5). A colouring of the edges of the cover graph of $\mathcal{P}_\Delta = \mathcal{P}_\Delta(D, c_l, c_u)$ with colours from $V \setminus \{v_0\}$ is naturally given by colouring a cover with the vertex of the corresponding vertex cut.

Lemma 3.7. *The colouring of the edges of the cover graph of \mathcal{P}_Δ with colours in $V \setminus \{v_0\}$ is a U -colouring.*

Proof. Let $x \in \mathcal{B}_\Delta(D, c_l, c_u)$ and suppose that pushing v in x is legal, *i.e.*, there is a covering coloured v leaving x . The other element x' of the covering pair is completely determined by x and v . This yields property U_1 .

For U_2 assume that u and v can both be pushed in x . We have to show that they can be pushed in either order. This clearly holds if the vertex cuts of u and v are disjoint. Suppose $(u, v) \in A$ and note that $c_l(a) < x(a) < c_u$. Now, after pushing u we still have $c_l(a) < x(a)$; thus we can still push at v . Conversely, pushing v preserves $x(a) < c_u$, *i.e.*, the push of u remains legal. □

A completely symmetric argument shows that the colouring is also a U-colouring for the reversed order, *i.e.*, an L-colouring. Theorem 2.13 implies that every connected component of \mathcal{P}_Δ is a distributive lattice. To complete the proof of Theorem 1.2 it only remains to show that \mathcal{P}_Δ is connected. This is shown in the last lemma of this section.

Lemma 3.8. *The order \mathcal{P}_Δ is connected.*

Proof. With Δ -bonds x and y consider $z = x - y$. Note that z is a **0**-bond because $\delta(C, z) = \delta(C, x) - \delta(C, y) = 0$ for all cycles C . Since **0**-bonds are just bonds, there is a unique expression of z as a linear combination of vertex cuts $S[v]$ with $v \neq v_0$; we write this as $z = \sum_v \lambda_v S[v]$. Based on the coefficients λ_v , define $T = \{v \in V : \lambda_v > 0\}$. Since x and y are different and could as well be exchanged, we may assume that $T \neq \emptyset$.

We claim that pushing T in y is legal, *i.e.*, $y' = \text{push}_T(y)$ is a Δ -bond. First note that $v_0 \notin T$. Now let $a = (v, w)$ be an arc with $v \in T$ and $w \in V \setminus T$, *i.e.*, $a \in S[T]^+$; from $(x - y)(a) = \lambda_v - \lambda_w > 0$ we obtain $x(a) > y(a)$, and hence $y(a) < c_u(a)$. For $a \in S[T]^-$ we obtain $x(a) < y(a)$, and hence $y(a) > c_l(a)$. Lemma 3.4 implies that y' is a Δ -bond.

Assuming that \mathcal{P}_Δ has several components we may choose x and y from different components such that $\sum_v \lambda_v$ is minimal, where $x - y = \sum_v \lambda_v S[v]$. Since y' is obtained from y by pushing the set T , it is in the same component as y . Note that

$$x - y' = \sum_{v \in V \setminus T} \lambda_v S[v] + \sum_{v \in T} (\lambda_v - 1) S[v].$$

Since $\sum_v \lambda_v - |T| < \sum_v \lambda_v$, this contradicts the choice of x and y . □

4. More applications of the ULD characterization

In the first three parts of this section we deal with special cases of Theorem 3.3. As a result we reprove known instances of distributive lattices from graphs. Section 4.4 links Δ -bonds, *i.e.*, special edge weightings, to potentials, *i.e.*, vertex weightings. In Section 4.5 we discuss the chip-firing game on directed graphs. A central and previously known result is that the states of this game carry the structure of a ULD lattice. We obtain this as a direct application of our characterization.

4.1. The lattice of c -orientations (Propp [23])

Given an orientation O of a graph $G = (V, E)$, we regard a cycle of an undirected graph as an edge sequence, rather than an edge set, *i.e.*, a cycle comes with the direction of its traversal. If C is a cycle in G , then denote by $c_O(C) := |C_O^+| - |C_O^-|$ the *circular flow-difference* of O around C , where C_O^+ is the set of forward arcs of C in O and C_O^- is the set of backward arcs.

Given a vector c , which assigns to every cycle C of G an integer $c(C)$, we call an orientation O of G with $c(C) = c_O(C)$ a *c -orientation*.

The main result in Propp’s article [23] is as follows.

Theorem 4.1. *Let $G = (V, A)$ be a graph and $c \in \mathbb{Z}^C$. The set of c -orientations of G carries the structure of a distributive lattice.*

Proof. Let $D = (V, A)$ be any orientation of G . Define $\Delta := \frac{1}{2}(c_D - c)$. We interpret $x \in \mathcal{B}_\Delta(D, \mathbf{0}, \mathbf{1})$ as the orientation $O(x)$ of G which arises from D by changing the orientation of $a \in A$ if $x(a) = 1$. For an arc set $A' \subseteq A$, let $x(A') := \sum_{a \in A'} x(a)$. We calculate

$$\begin{aligned} c_{O(x)}(C) &= |C_{O(x)}^+| - |C_{O(x)}^-| \\ &= |C_D^+| - x(C_D^+) + x(C_D^-) - (|C_D^-| - x(C_D^-) + x(C_D^+)) \\ &= |C_D^+| - |C_D^-| - 2(x(C_D^+) - x(C_D^-)) \\ &= c_D(C) - 2\delta(x, C) = c_D(C) - 2\Delta_C \\ &= c. \end{aligned}$$

This shows that c -orientations of G correspond bijectively to Δ -bonds in $\mathcal{B}_\Delta(D, \mathbf{0}, \mathbf{1})$. By Theorem 3.3 we obtain a distributive lattice structure on the set of c -orientations of G . \square

From duality of planar graphs and the above theorem, Propp derives the following two corollaries.

- The set of d -factors of a plane bipartite graph can be enhanced by a distributive lattice structure.
- The set of spanning trees of a plane graph can be enhanced by a distributive lattice structure.

4.2. The lattice of flow in planar graphs (Khuller, Naor and Klein [17])

Consider a planar digraph $D = (V, A)$, with each arc a having an integer *lower* and *upper* bound on its capacity, denoted $c_\ell(a)$ and $c_u(a)$. For a function $f : A \rightarrow \mathbb{Z}$, call

$$\omega(v, f) := \sum_{a \in \text{in}(v)} f(a) - \sum_{a \in \text{out}(v)} f(a)$$

the *excess* at v . Given a vector $\Omega \in \mathbb{N}^V$, call f a Ω -flow if $c_\ell(a) \leq f(a) \leq c_u(a)$ for all a and $\Omega_v = \omega(v, f)$ for all $v \in V$. Let $\mathcal{F}_\Omega(D, c_\ell, c_u)$ denote the set of Ω -flows.

Theorem 4.2. *If D is a planar digraph then $\mathcal{F}_\Omega(D, c_\ell, c_u)$ carries the structure of a distributive lattice.*

Proof. Given a crossing-free plane embedding of D , we look at the planar dual digraph D^* , which is an orientation of the planar dual G^* of the underlying undirected graph G of D . Let v be a vertex of G^* corresponding to a facial cycle C of the embedding of D . Orient an edge incident to v as an outgoing arc of v if the dual arc is forward when traversing C in clockwise direction. Given values on the arcs of D , we simply transfer them to the corresponding arcs of D^* .

Since the excess at a vertex of D dualizes to the circular flow difference around the corresponding facial cycle of D^* , we have a correspondence between $\mathcal{F}_\Omega(D, c_\ell, c_u)$ and $\mathcal{B}_\Omega(D^*, c_\ell, c_u)$. This yields the distributive lattice structure on Ω -flows of planar graphs. \square

In analogy to the case of Δ -bonds we can assume the data (D, c_ℓ, c_u, Ω) to be reduced. Now the dual operation to vertex pushes is to augment the flow around facial cycles. A natural candidate for the *forbidden facial cycle* is the unbounded face of the planar embedding. By flow-augmentation at the remaining facial cycles we can construct the cover graph of a distributive lattice on $\mathcal{F}_\Omega(D, c_\ell, c_u)$.

Khuller, Naor and Klein [17] in fact only consider the special case of Theorem 4.2 where $\Omega = \mathbf{0}$; these flows without excess are called *circulations*. We restate their result as follows.

Theorem 4.3. *Let D be a planar digraph with upper and lower arc capacities c_l and c_u . The set of circulations of D within c_l and c_u carries the structure of a distributive lattice.*

4.3. The lattice of α -orientations in planar graphs (Felsner [13])

Consider a plane graph $G = (V, E)$. Given a map $\alpha : V \rightarrow \mathbb{N}$, an orientation X of the edges of G is called an α -orientation if α records the out-degrees of all vertices, *i.e.*, $\text{outdeg}_X(v) = \alpha(v)$ for all $v \in V$.

The main result in [13] is as follows.

Theorem 4.4. *Given a planar graph and a mapping $\alpha : V \rightarrow \mathbb{N}$, the set of α -orientations of G carries the structure of a distributive lattice.*

Proof. In analogy to the proof of Theorem 4.1, where c -orientations were interpreted as elements of $\mathcal{B}_\Delta(C, \mathbf{0}, \mathbf{1})$, we can view α -orientations as elements of $\mathcal{F}_\Omega(D, \mathbf{0}, \mathbf{1})$ for some orientation D of G and Ω depending on α . Application of Theorem 4.2 yields the distributive lattice structure on the α -orientations. \square

In the perspective of the present paper, α -orientations appear to be a special case of the preceding constructions. Nevertheless they are quite general objects. Special instances of α -orientations on plane graphs yield lattice structures on:

- Eulerian orientations of a plane graph,
- spanning trees and d -factors of a plane graph,
- Schnyder woods of a 3-connected plane graph.

It is natural to ask for the structure of α -orientations of graphs embedded on some surface. Propp [23] comments that to move between the d -factors in toroidal graphs it is necessary to operate on non-contractible cycles. The second author's diploma thesis [18] investigates α -orientations on surfaces and generalizations of α -orientations.

4.4. Potentials

For a last and fundamental example set $\Delta = \mathbf{0}$. Fix $v_0 \in V$ and let $x(A') := \sum_{a \in A'} x(a)$, where $A' \subseteq A$. Now $x \in \mathcal{B}_0(D, c_\ell, c_u)$ means that the potential mapping $\pi_x(v) := x(p)$ for a (v_0, v) -path p is well defined, *i.e.*, independent of the choice of p . In particular, for every $v_0 \in V$ the set $\mathcal{B}_0(D, c_\ell, c_u)$ is in bijection with the set of *feasible vertex-potentials*:

$$\Pi_{v_0}(D, c_\ell, c_u) := \{\pi \in \mathbb{N}^V \mid \pi(v_0) = 0 \text{ and } c_\ell(a) \leq \pi(w) - \pi(w) \leq c_u(a) \text{ for all } a \in A\}.$$

Theorem 3.3 yields that $\Pi_{v_0}(D, c_\ell, c_u)$ carries the structure of a distributive lattice. Indeed, this can be obtained more easily, since $\Pi_{v_0}(D, c_\ell, c_u)$ is a suborder of the distributive lattice \mathbb{N}^V where the order is given by dominance and the lattice operations are max and min. The simple observation that, for $\pi, \pi' \in \Pi_{v_0}(D, c_\ell, c_u)$, their componentwise minimum and maximum are also in $\Pi_{v_0}(D, c_\ell, c_u)$, yields the distributive lattice structure on feasible vertex potentials.

The proof given in [13], [17] and [23] for the respective lattice structures are all based on the construction of potentials corresponding to the objects they investigate. In a forthcoming paper [14] we exploit the potential approach. In this paper we will deal with a class of *distributive polytopes*, *D-polytopes* for short, and we generalize the notion of Δ -bonds to *generalized bonds*. These are in a certain sense the most general structure on graphs that form a distributive lattice.

4.5. Chip-firing games

ULDs can be represented as *antimatroids with repetition* [7] or *locally free, permutable, left-hereditary languages* [6]. Theorem 1.1 gives an easy way to obtain such representations. In the following we apply Theorem 1.1 to the class of *chip-firing games*. As in [5] we consider chip-firing games on directed graphs. For the undirected case see [6]. In general the ULDs coming from chip-firing games are not distributive.

Let $D = (V, A)$ be a directed graph with a map $\sigma_0 : V \rightarrow \mathbb{N}$ called a *chip arrangement*. The number $\sigma_0(v)$ records the *number of chips* on vertex v in σ_0 . Given a chip arrangement σ a vertex v can be *fired* if it contains more chips than its out-degree, *i.e.*, $\sigma(v) \geq \text{outdeg}(v)$. Firing v consists of sending a chip along each of the out-going arcs of v to their respective end-vertices. The new chip arrangement is called σ^v . Define a directed graph $\text{CFG}(D, \sigma_0)$ on the set of chip arrangements which can be obtained from σ_0 by a firing sequence. The arc (σ, σ^v) is naturally coloured with $v \in V$. By definition of $\text{CFG}(D, \sigma_0)$, every vertex σ lies on a directed (σ_0, σ) -path. If $\text{CFG}(D, \sigma_0)$ is also acyclic the chip-firing game is called *finite*. We obtain the following well-known result.

Proposition 4.5. *The states of a finite chip-firing game carry the structure of a ULD.*

Proof. Since $\text{CFG}(D, \sigma_0)$ is acyclic and each vertex lies on a directed (σ_0, σ) -path, $\text{CFG}(D, \sigma_0)$ is connected and has a unique sink σ_0 . It only remains to show that the natural colouring of $\text{CFG}(D, \sigma_0)$ is a U-colouring.

Clearly we have property U_1 . To prove property U_2 , let $\sigma \in \text{CFG}(D, \sigma_0)$ and $v, w \in V$ ready for firing in σ . Since the firing of v can only increase the number of chips on w , after firing v vertex w can still be fired. This means that colouring the arcs of $\text{CFG}(D, \sigma_0)$ by the vertices that have been fired leads to a U-colouring. Hence, by Theorem 1.1, the digraph $\text{CFG}(D, \sigma_0)$ is the cover graph of a ULD. □

The ULD-properties imply that in a finite game there is a unique chip arrangement σ^* where the game starting in σ_0 ends, and that all firing sequences from a given chip arrangement σ to the maximum σ^* fire the same multiset of vertices.

Observe that the property which makes a U-poset a ULD, namely the existence of a global minimum, is rather artificially achieved in the case of chip-firing games. Instead of considering only the arrangements that can be reached from a starting chip arrangement σ_0 by *upwards* transformations, we can endow the structure by the inverse operation of firing and call it *co-firing*. Co-firing a vertex v means sending one chip along all the arcs (v, w) from w to v . Note that this requires that every out-neighbour of v owns a chip. We define the *complete chip-firing game* $\text{CCFG}(D, \sigma_0)$ to be the digraph on all the chip arrangements that can be reached from σ_0 by firing and co-firing vertices. Under the assumption that $\text{CFG}(D, \sigma_0)$ is acyclic, the obvious colouring

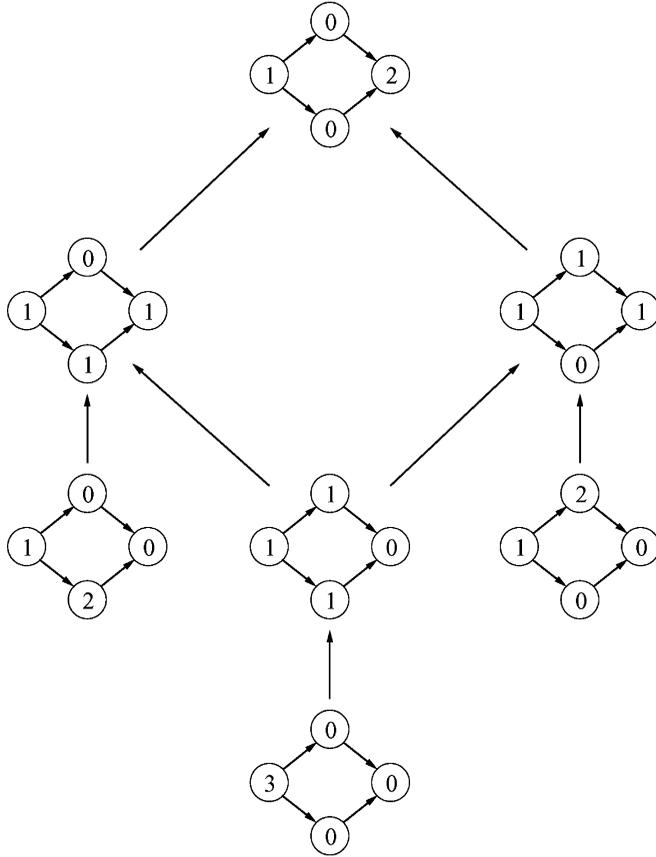


Figure 5. A U-poset induced by a complete chip-firing game.

of the edges makes $CCFG(D, \sigma_0)$ a U-poset. This U-poset is not a lattice, but all its principal upsets (filters) correspond to ULDs obtained from chip-firing games in the original sense; see, e.g., Figure 5.

It could be worthwhile to investigate the U-posets coming from complete chip-firing games. As a first step we provide the following proposition related to Remark 2. Note that the property stated in the proposition does not hold for the U-poset of Figure 4.

Proposition 4.6. *If (P, \leq) is a U-poset induced by a CCFG, then for every $s \in P$ there is a unique inclusion-minimal set M_s of meet-irreducibles such that $\bigwedge M_s = s$.*

Proof. By Remark 2 we know that for every $s \in P$ there is a unique inclusion-minimal set M_s of meet-irreducibles such that $s \in \bigwedge M_s$. So suppose there are distinct $s, t \in \bigwedge M_s$, i.e., $M_s = M_t$. Partition $M_s = U \cup V$ and let $u = \bigwedge U$ and $v = \bigwedge V$ and $w = u \vee v$. The relation between these elements is as in Figure 4.

For two states $i < j$ in $CCFG(D, \sigma_0)$, let $p_{i,j}$ be a directed (i, j) -path and $c(i, j)$ be the colourset of $p_{i,j}$. Recall from Section 2 that $c(i, j)$ is indeed independent of the choice of $p_{i,j}$. Since

$s = u \wedge v$ and $w = u \vee v$ we have $c(s, u) = c(v, w)$ and $c(s, v) = c(u, w)$. Similarly, for t we obtain $c(t, u) = c(v, w)$ and $c(t, v) = c(u, w)$. This yields $c(s, u) = c(t, u)$ and $c(s, v) = c(t, v)$. But if a chip-configuration $u \in \text{CCFG}(D, \sigma_0)$ can be obtained from s and t by firing the same multiset of vertices, then $s = t$. This is $\bigwedge M_s = s$. \square

5. Concluding remarks

Our characterization of ULDs originates from a characterization of matrices whose *flip-flop* poset generates a distributive lattice in [18]. It turned out that this tool yields handy proofs for the known distributive lattices from graphs. The extraction of the data that were necessary for the proof led to Δ -bonds. While preparing the present paper we observed that every Δ -bond lattice is isomorphic to a $\mathbf{0}$ -bond lattice, and hence the lattice structure can be nicely proved via potentials. This led us to the notion of a D-polytope, *i.e.*, of a polytope P such that, with points x and y , $\max(x, y)$ and $\min(x, y)$ are also in P . We can characterize the bounding hyperplanes of D-polytopes. This allows us to associate a weighted digraph with every D-polytope such that the *generalized bonds* of the weighted digraph form a distributive lattice. This is the topic of the forthcoming paper [14].

There are two types of problem closely related to this article where we would like to see progress.

- Lattices of Δ -bonds depend on the choice of a forbidden vertex $v_0 \in V$. Choosing another forbidden vertex v_1 yields a different lattice on the same set of objects. Is there an easy description of the transformation $\mathcal{P}_\Delta(v_0) \rightarrow \mathcal{P}_\Delta(v_1)$?
- The generation of a random element from a distributive lattice is a nice application for *coupling from the past* (*cf.* Propp and Wilson [25]). The challenge is to find good estimates for the mixing time: see Propp [24]. What if the lattice is a Δ -bond lattice?

Acknowledgements

Thanks to Éric Fusy for fruitful discussions during his stay in Berlin and to Edwin O’Shea for his help with the exposition.

References

[1] Ardila, F. and Maneva, E. (2007) Pruning processes and a new characterization of convex geometries. *Discrete Math.* **309** 3083–3091. [arXiv:0706.3750v4 \[math.CO\]](https://arxiv.org/abs/0706.3750v4).

[2] Avann, S. P. (1968) Locally atomic upper locally distributive lattices. *Math. Ann.* **175** 320–336.

[3] Birkhoff, G. and Bennett, M. K. (1985) The convexity lattice of a poset. *Order* **2** 223–242.

[4] Bjorner, A. (1980) Shellable and Cohen–Macaulay partially ordered sets. *Trans. Amer. Math. Soc.* **260** 159–183.

[5] Bjorner, A. and Lovasz, L. (1992) Chip-firing games on directed graphs. *J. Algebraic Combin.* **1** 305–328.

[6] Bjorner, A., Lovasz, L. and Shor, P. W. (1991) Chip-firing games on graphs. *European J. Combin.* **12** 283–291.

[7] Bjorner, A. and Ziegler, G. M. (1992) Introduction to greedoids. In *Matroid Applications*, Vol. 40 of *Encyclopedia of Mathematics and its Applications*, Cambridge University Press, Cambridge, pp. 284–357.

- [8] Boulaye, G. (1967) Sous-arbres et homomorphismes à classes connexes dans un arbre. In *Theory of Graphs International Symposium*, pp. 47–50.
- [9] Dilworth, R. P. (1940) Lattices with unique irreducible decompositions. *Ann. of Math.* **41** 771–777.
- [10] Dilworth, R. P. and Crawley, P. (1960) Decomposition theory for lattices without chain conditions. *Trans. Amer. Math. Soc.* **96** 1–22.
- [11] Edelman, P. H. (1980) Meet-distributive lattices and the antiexchange closure. *Alg. Univ.* **10** 290–299.
- [12] Edelman, P. H. and Jamison, R. E. (1985) The theory of convex geometries. *Geom. Dedicata* **19** 247–270.
- [13] Felsner, S. (2004) Lattice structures from planar graphs. *Electron. J. Combin.* **11** #15.
- [14] Felsner, S. and Knauer, K. B. (2008) Distributive lattices, polyhedra, and generalized flow. [arXiv:0811.1541v1](https://arxiv.org/abs/0811.1541v1) [math.CO].
- [15] Jamison, R. E. (1970) A development of axiomatic convexity. *Clemson University Math.* **48** 15–20.
- [16] Jamison, R. E. (1980) Copoints in antimatroids. *Congr. Numer.* **29** 534–544.
- [17] Khuller, S., Naor, J. and Klein, P. (1993) The lattice structure of flow in planar graphs. *SIAM J. Discrete Math.* **6** 477–490.
- [18] Knauer, K. B. (2007) Partial orders on orientations via cycle flips. <http://www.math.tu-berlin.de/~knauer/diplom.pdf>.
- [19] Korte, B. and Lovasz, L. (1984) Shelling structures, convexity and a happy end. In *Graph Theory and Combinatorics* (Cambridge 1983), Academic Press, London, pp. 219–232.
- [20] Monjardet, B. (1990) The consequences of Dilworth’s work on lattices with unique irreducible decompositions. In *The Dilworth Theorems*, Birkhäuser, pp. 192–201.
- [21] Nakamura, M. (2003) Excluded-minor characterizations of antimatroids arisen from posets and graph searches. *Discrete Appl. Math.* **129** 487–498.
- [22] Pfaltz, J. L. (1971) Convexity in directed graphs. *J. Combin. Theory Ser. B* **10** 143–152.
- [23] Propp, J. (1993) Lattice structure for orientations of graphs. [arXiv:math/0209005v1](https://arxiv.org/abs/math/0209005v1) [math.CO].
- [24] Propp, J. (1997) Generating random elements of finite distributive lattices. *Electron. J. Combin.* **4** #15.
- [25] Propp, J. and Wilson, D. B. (1996) Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Struct. Alg.* **9** 223–252.
- [26] Stern, M. (1999) *Semimodular Lattices*, Vol. 73 of *Encyclopedia of Mathematics and its Applications*, Cambridge University Press, Cambridge.