

Independence posets

HUGH THOMAS* AND NATHAN WILLIAMS†

Let G be an acyclic directed graph. For each vertex $g \in G$, we define an involution on the independent sets of G . We call these involutions flips, and use them to define a new partial order on independent sets of G .

Trim lattices generalize distributive lattices by removing the graded hypothesis: a graded trim lattice is a distributive lattice, and every distributive lattice is trim. Our independence posets are a further generalization of distributive lattices, eliminating also the lattice requirement: an independence poset that is a lattice is always a trim lattice, and every trim lattice is the independence poset for a unique (up to isomorphism) acyclic directed graph G . We characterize when an independence poset is a lattice with a graph-theoretic condition on G .

We generalize the definition of rowmotion from distributive lattices to independence posets, and we show it can be computed in three different ways. We also relate our constructions to torsion classes, semibricks, and 2-simpleminded collections arising in the representation theory of certain finite-dimensional directed algebras.

1. Introduction

In this paper, we always take G to be a finite acyclic directed graph. The transitive closure of G defines a poset, which we refer to as G -order. Our convention is that $g_1 \geq g_2$ in G -order if and only if there is a directed path in G from g_1 to g_2 ; when we compare vertices of G , we will always mean a comparison in G -order. We write \simeq for an isomorphism of posets.

1.1. Independent sets and tight orthogonal pairs

Recall that an *independent set* $\mathcal{A} \subseteq G$ is a set of pairwise non-adjacent vertices of G . As we now explain, the orientation provided by G allows us to

arXiv: [1805.00815](https://arxiv.org/abs/1805.00815)

*Partially supported by the Canada Research Chairs program and an NSERC Discovery Grant.

†Partially supported by a Simons Foundation Collaboration Grant.

complete an independent set to a pair of independent sets, either of which determines the other.

Definition 1.1. A pair $(\mathcal{D}, \mathcal{U})$ of independent sets of G is called *orthogonal* if there is no edge in G from an element of \mathcal{D} to an element of \mathcal{U} . An orthogonal pair of independent sets $(\mathcal{D}, \mathcal{U})$ is called *tight* if whenever any element of \mathcal{D} is increased (removed and replaced by a larger element with respect to G -order) or any element of \mathcal{U} is decreased, or a new element is added to either \mathcal{D} or \mathcal{U} , then the result is no longer an orthogonal pair of independent sets. We abbreviate **tight orthogonal pair** by **top**, and we write $\text{top}(G)$ for the set of all tops of G .

Some examples are given in Figure 1.




Figure 1: Eight pairs of independent sets $(\mathcal{D}, \mathcal{U})$ for two different orientations of a path graph. The blue (dark gray) vertices correspond to the elements of \mathcal{D} , while the orange (light gray) vertices correspond to \mathcal{U} .

An independent set can be completed to a tight orthogonal pair in exactly two ways (see Algorithms 1 and 2, and Figure 4).

Theorem 1.2. Let \mathcal{I} be an independent set of a directed acyclic graph G . Then there exists a unique $(\mathcal{I}, \mathcal{U}) \in \text{top}(G)$ and a unique $(\mathcal{D}, \mathcal{I}) \in \text{top}(G)$.

1.2. Flips and the independence poset

Fix ℓ a linear extension of G -order and ℓ' a reverse linear extension of G -order. Note that by our conventions, a linear extension is a linear order such that if there is an edge $g_1 \rightarrow g_2$, then g_2 precedes g_1 in the linear extension.

Definition 1.3. The *flip* of $(\mathcal{D}, \mathcal{U}) \in \text{top}(G)$ at an element $g \in G$ is the tight orthogonal pair $\text{flip}_g(\mathcal{D}, \mathcal{U})$ defined as follows (see Figure 2 for an example): if $g \notin \mathcal{D}$ and $g \notin \mathcal{U}$, the flip does nothing. Otherwise, preserve all elements of \mathcal{D} that are not less than g and all elements of \mathcal{U} that are not greater than g (and delete all other elements); after switching the set to which g belongs, then greedily complete \mathcal{D} and \mathcal{U} to a tight orthogonal pair in the orders ℓ' and ℓ , respectively.




Figure 2: A flip on a top $(\mathcal{D}, \mathcal{U})$ in the 7×7 grid oriented from top left to bottom right. As in Figure 1, the blue vertices correspond to the elements of \mathcal{D} , while the orange vertices correspond to the elements of \mathcal{U} . Flipping at the vertex g changes its color, and divides the grid into 5 connected regions (delineated by the dotted lines): the blue vertices not less than g (i.e., not in the bottom right) and the orange vertices not greater than g (i.e., not in the top left) are preserved by the flip. The orange vertices in the top left are filled in greedily from bottom right to top left; the blue vertices in the bottom right are filled in greedily from top left to bottom right.

Pseudocode for Definition 1.3 is given in Algorithm 3. Proposition 3.1 proves that the algorithm produces a tight orthogonal pair, while Lemma 3.2 proves that flips are involutions.

Definition 1.4. We define the *independence poset* on $\text{top}(G)$ as the reflexive and transitive closure of the relations $(\mathcal{D}, \mathcal{U}) \prec (\mathcal{D}', \mathcal{U}')$ if there is some $g \in \mathcal{U}$ such that $\text{flip}_g(\mathcal{D}, \mathcal{U}) = (\mathcal{D}', \mathcal{U}')$. (Lemma 3.3 proves that this really does define a poset, and that these relations are exactly its covers.)

We denote this poset by $\text{top}(G)$. By construction, $\text{top}(G)$ is connected and has a minimum and a maximum element. Figure 3 gives some examples of independence posets on various orientations of a path of length four, while Figure 6 realizes the Tamari lattice on 14 elements as an independence poset. As we summarize below, trim lattices are special cases of independence posets, and so the class of independence posets includes all distributive lattices, Tamari lattices, Cambrian lattices, Fuss-Cambrian lattices, and torsion pairs of tilted finite type hereditary Artin algebras.




Figure 3: Independence posets for four orientations of a path of length 4. Each poset has eight elements (the tops drawn in blue and orange as in Figures 1 and 2), corresponding to the eight independent sets in the underlying undirected graph. The top left poset is not a lattice, the bottom left poset is a distributive lattice, and both posets on the right are trim lattices.

Theorem 1.5. *Fix a linear extension ℓ of G -order. Flipping only in increasing order of ℓ gives a tree structure on the independent sets of G .*

1.3. Independence lattices are trim

A *trim lattice* is an extremal left-modular lattice [Tho06]. Trim lattices were introduced to serve as analogues of distributive lattices without the graded hypothesis: a graded trim lattice is a distributive lattice, and every distributive lattice is trim.

Our independence posets further generalize distributive lattices by removing the lattice requirement: an independence poset that is a lattice is always a trim lattice, and every trim lattice can be realized as an independence poset for a unique (up to isomorphism) acyclic directed graph G . In other words, the common intersection of lattices and independence posets are exactly the trim lattices.

Theorem 1.6. *If $\text{top}(G)$ is a lattice, then it is a trim lattice.*

Following Markowsky [Mar75, Mar92, TW17], a **maximal orthogonal pair** (or mop) in an acyclic directed graph G is a pair of sets (X, Y) such that no edges run from X to Y , and such that X and Y are both maximal with respect to this condition. Markowsky's generalization of Birkhoff's fundamental theorem of finite distributive lattices states that any extremal lattice has a representation $\mathcal{L}(G)$ as the lattice of maximal orthogonal pairs of a unique (up to isomorphism) acyclic directed graph G .

Theorem 1.7. *If $\mathcal{L}(G)$ is trim, then $\mathcal{L}(G) \simeq \text{top}(G)$.*

Moreover, if $\text{top}(G)$ is a lattice, then $\text{top}(G) \simeq \mathcal{L}(G)$. We provide explicit bijections between tops and mops in Section 4.4. The common lattice of tops and mops offer different advantages. Cover relations $x < y \in \text{top}(G)$ are easy to compute using tops (by flips), but harder to see using mops. Similarly, relations $x < y \in \mathcal{L}(G)$ are easy to compute using mops (by inclusion), but harder to see using tops.

Theorem 4.11 characterizes when $\text{top}(G)$ is a lattice—or, equivalently, when $\mathcal{L}(G)$ is trim—in terms of graph-theoretic properties of G . We discuss some properties of independence posets that are not lattices in Section 4.6.

1.4. Toggles

By Theorem 1.2, the number of elements of $\text{top}(G)$ is equal to the number of independent sets in the *undirected* graph G . For g a minimal or maximal element of G , there is a natural *toggle* operation tog_g (similar to quiver mutation) that reverses every edge incident to g ; this operation induces a bijection between $\text{top}(G)$ and $\text{top}(\text{tog}_g(G))$ whose effect essentially inter-

changes the relative order of a decomposition of $\text{top}(G)$ into two intervals using the element g .

By limiting ourselves to only toggling at maximal elements (or at minimal elements) and only keeping track of the sets \mathcal{U} (or \mathcal{D}), this bijection can be computed cleanly at the level of independent sets (see Equation (3) and Theorem 5.1). Some examples of toggles on an orientation of the path of length three are given in Figure 8. See also [SW12, Str16].

1.5. Rowmotion

Since both components of a top are independent sets, and each independent set can be completed to a top in two ways, it is natural to define *rowmotion* by sending one completion to the other:

$$(1) \quad \text{row}(\mathcal{D}, \mathcal{U}) := \text{the unique } (\mathcal{D}', \mathcal{U}') \in \text{top}(G) \text{ with } \mathcal{D} = \mathcal{U}'.$$

It turns out that there are two other equally natural (but slower) ways to compute rowmotion. The distinction between these two slower methods was not apparent when rowmotion had been studied at the level of distributive lattices [Fon93, CF95, SW12], but our generalized setting of independence posets reveals their differences: one method computes rowmotion as a composition of flips within a *fixed* independence poset (*rowmotion in slow motion*), while the second relies on a sequence of toggles and the corresponding bijections between independence posets for different orientations of the same underlying undirected graph (*rowmotion by deformation*).

Theorem 1.8. *Let G be a directed acyclic graph. Then rowmotion can be computed in slow motion and by deformation—that is,*

$$\text{row} = \prod_{g \in \ell} \text{flip}_g = \prod_{g \in \ell'} \text{tog}_g$$

for any linear extension ℓ and reverse linear extension ℓ' of G -order.

1.6. Representation theory

We conclude with some applications to representation theory. Let k be a field, and A a finite-dimensional k -algebra such that the module category $\text{mod } A$ has no cycles. Define a directed graph G with vertices indexed by the indecomposable A -modules and an arrow from M to N if and only if $\text{Hom}(M, N) \neq 0$. By our assumption that there are no cycles in $\text{mod } A$, the graph G is acyclic—but not all acyclic directed graphs G arise in this way.

A *torsion class* in $\text{mod } A$ is a full additive subcategory of $\text{mod } A$ closed under extensions and quotients. A module is called a *brick* if its endomorphism ring is a division algebra. A collection of bricks is called a *semibrick* if there are no morphisms between two non-isomorphic bricks in the collection. Finally, *2-simple-minded collections* (defined in Section 7.4) are certain collections of objects in the derived category of A in bijection with torsion classes [Asa16].

Theorem 1.9. *If A is representation finite and $\text{mod } A$ has no cycles, then maximal orthogonal pairs correspond to torsion pairs, independent sets correspond to semibricks, and tight orthogonal pairs correspond to 2-simple-minded collections.*

1.7. Organization of the paper

In Section 2 we complete independent sets to tight orthogonal pairs with Algorithms 1 and 2, proving Theorem 1.2. In Section 3.1, we define flips with Algorithm 3, and prove they are well-defined in Proposition 3.1. We study independence posets in Section 3.2, and present a useful recursion on tops in Section 3.3. We relate independence posets to trim lattices in Section 4, relating tight orthogonal pairs to maximal orthogonal pairs and proving Theorems 1.6 and 1.7. In Section 5, we define toggles on independence posets; we then study rowmotion in Section 6, proving Theorem 1.8. We conclude with connections to representation theory in Section 7.

2. Tight orthogonal pairs

Definition 2.1. Let $g \in G$ and define G_g to be the directed graph obtained by deleting the vertex g from G (along with all edges to g), and G_g° the directed graph obtained by deleting all vertices and edges adjacent to g in G (along with g itself).

Theorem 1.2. *Let \mathcal{F} be an independent set of an acyclic directed graph G . Then there exists a unique $(\mathcal{F}, \mathcal{U}) \in \text{top}(G)$ and a unique $(\mathcal{D}, \mathcal{F}) \in \text{top}(G)$.*

Proof. Let \mathcal{F} be an independent set. We show that Algorithm 1 produces an element $(\mathcal{D}, \mathcal{F}) \in \text{top}(G)$ (this algorithm is illustrated in Figure 4). By construction, the output of Algorithm 1 is an orthogonal pair of independent sets. We claim it is tight.

Suppose the output is not tight. Then at least one of the following holds:

- there is some element $g \in G$ that could be added to \mathcal{D} to still have an orthogonal pair of independent sets,

```

Input: An acyclic directed graph  $G$  and an independent set  $\mathcal{F}$ .
Output: An element  $(\mathcal{D}, \mathcal{F}) \in \text{top}(G)$ .
set:  $\mathcal{D} = \{\}$ 
for  $k$  in  $\ell'$  do
  if  $\left\{ \begin{array}{l} k \notin \mathcal{F} \\ i \rightarrow k \notin G \text{ for } i \in \mathcal{D} \\ k \rightarrow i \notin G \text{ for } i \in \mathcal{F} \end{array} \right\}$  then add  $k$  to  $\mathcal{D}$ 
end
return  $(\mathcal{D}, \mathcal{F})$ 

```

Algorithm 1: The greedy construction of the unique $(\mathcal{D}, \mathcal{F}) \in \text{top}(G)$ using any reverse linear extension ℓ' of G -order, given an independent set \mathcal{F} . See Figure 4 for an example.

```

Input: An acyclic directed graph  $G$  and an independent set  $\mathcal{F}$ .
Output: An element  $(\mathcal{F}, \mathcal{U}) \in \text{top}(G)$ .
set:  $\mathcal{U} = \{\}$ 
for  $k$  in  $\ell$  do
  if  $\left\{ \begin{array}{l} k \notin \mathcal{F} \\ i \rightarrow k \notin G \text{ for } i \in \mathcal{F} \\ k \rightarrow i \notin G \text{ for } i \in \mathcal{U} \end{array} \right\}$  then add  $k$  to  $\mathcal{U}$ 
end
return  $(\mathcal{F}, \mathcal{U})$ 

```

Algorithm 2: The greedy construction of the unique $(\mathcal{F}, \mathcal{U}) \in \text{top}(G)$ using any linear extension ℓ of G -order, given an independent set \mathcal{F} . See Figure 4 for an example.

- there is some element $g' \in \mathcal{D}$ that could be increased to $g \in G$ with respect to G -order,
- there is some element $g \in G$ that could be added into \mathcal{F} , or
- there is some element $g' \in \mathcal{F}$ that could be decreased to $g \in G$.

Take a maximal g among all such elements. One verifies that Algorithm 1 would have added it in, which is a contradiction. Similar reasoning shows that Algorithm 2 produces an element $(\mathcal{F}, \mathcal{U}) \in \text{top}(G)$.

We now show by induction on $|G|$ that given \mathcal{F} , the tight orthogonal pair $(\mathcal{D}, \mathcal{F})$ is unique. The base case for $|G| = 1$ is trivial. We now suppose $(\mathcal{D}, \mathcal{F})$ and $(\mathcal{D}', \mathcal{F})$ are two different tight orthogonal pairs. Let g be minimal in G .

If $g \notin \mathcal{F}$, the restriction of $(\mathcal{D}, \mathcal{F})$ and $(\mathcal{D}', \mathcal{F})$ to G_g are tight orthogonal pairs of G_g ; by induction, they must coincide except possibly at g . But since $\mathcal{D} \neq \mathcal{D}'$ then either $\mathcal{D} \subset \mathcal{D}'$ or $\mathcal{D}' \subset \mathcal{D}$, which contradicts tightness.




Figure 4: The progression $(a) \mapsto (b) \mapsto (c)$ is an illustration of Algorithm 1, which greedily adds elements to the independent set (a) in a reverse linear extension of G -order (from the top left to the bottom right). The progression $(d) \mapsto (e) \mapsto (f)$ illustrates Algorithm 2, which greedily adds elements to the independent set (d) in a linear extension of G -order (from the bottom right to the top left). Rowmotion, defined by Equation (1), sends the tight orthogonal pair (f) on the bottom to the tight orthogonal pair (c) .

Otherwise, $g \in \mathcal{F}$. The restriction of $(\mathcal{D}, \mathcal{F})$ and $(\mathcal{D}', \mathcal{F})$ to G_g° are now tight orthogonal pairs of G_g° and coincide. But by definition, there can be no elements of either \mathcal{D} or \mathcal{D}' adjacent to g and so $(\mathcal{D}, \mathcal{F})$ and $(\mathcal{D}', \mathcal{F})$ coincide.

The argument that the tight orthogonal pair $(\mathcal{F}, \mathcal{U})$ produced by Algorithm 2 is unique is similar, instead letting g be maximal in G . \square

3. Flips and independence posets

3.1. Flips on tight orthogonal pairs

We will define a poset structure on the tight orthogonal pairs of G by specifying the cover relations. To this end, Definition 1.3 and Algorithm 3 define a *flip* of a tight orthogonal pair $(\mathcal{D}, \mathcal{U})$ at an element $g \in G$, written $\text{flip}_g(\mathcal{D}, \mathcal{U})$; the flip moves *up* in the poset if $g \in \mathcal{U}$ and moves *down* in the poset if $g \in \mathcal{D}$ (and does nothing otherwise). Figure 2 illustrates a flip on a tight orthogonal pair in an orientation of $[7] \times [7]$. We first prove that the image of a flip is again a tight orthogonal pair.

Proposition 3.1. *Let g be an element of an acyclic directed graph G . Then $\text{flip}_g(\mathcal{D}, \mathcal{U}) \in \text{top}(G)$.*

Proof. The statement follows from the restriction of Theorem 1.2 to the elements of G not less than g and to the elements not greater than g . \square

```

Input:  $(\mathcal{D}, \mathcal{U}) \in \text{top}(G)$  and  $g \in G$ .
Output:  $(\mathcal{D}', \mathcal{U}') \in \text{top}(G)$ .
set:  $\mathcal{D}' = \{k : k \in \mathcal{D} \text{ and } k \not\geq g\}$  and  $\mathcal{U}' = \{k : k \in \mathcal{U} \text{ and } k \not\geq g\}$ 
if  $g \notin \mathcal{U}$  and  $g \notin \mathcal{D}$  then return  $(\mathcal{D}, \mathcal{U})$  else if  $g \in \mathcal{U}$  then
 $\mathcal{D}' = \mathcal{D}' \cup \{g\}$  else if  $g \in \mathcal{D}$  then  $\mathcal{U}' = \mathcal{U}' \cup \{g\}$  for  $k$  in  $\ell'$  do
    if  $\left\{ \begin{array}{l} k \not\geq g, k \notin \mathcal{U} \\ k \rightarrow i \notin G \text{ for } i \in \mathcal{U} \\ i \rightarrow k \notin G \text{ for } i \in \mathcal{D} \end{array} \right\}$  then add  $k$  to  $\mathcal{D}'$ 
end
for  $k$  in  $\ell$  do
    if  $\left\{ \begin{array}{l} k \not\geq g, k \notin \mathcal{D} \\ i \rightarrow k \notin G \text{ for } i \in \mathcal{D} \\ k \rightarrow i \notin G \text{ for } i \in \mathcal{U} \end{array} \right\}$  then add  $k$  to  $\mathcal{U}'$ 
end
return  $(\mathcal{D}', \mathcal{U}')$ 

```

Algorithm 3: The definition of a flip of $(\mathcal{D}, \mathcal{U}) \in \text{top}(G)$ at an element $g \in G$, written $\text{flip}_g(\mathcal{D}, \mathcal{U})$. As usual, ℓ is a linear extension of G -order, while ℓ' is a reverse linear extension.

Lemma 3.2. *Let g be an element of an acyclic directed graph G . Then*

$$\text{flip}_g^2(\mathcal{D}, \mathcal{U}) = (\mathcal{D}, \mathcal{U}).$$

If h is incomparable with g in G -order, then $\text{flip}_g \circ \text{flip}_h = \text{flip}_h \circ \text{flip}_g$.

Proof. For the first statement, a flip preserves the elements of \mathcal{D} that are not less than g , and the elements of \mathcal{U} that are not greater than g . Since g has now returned to its original set after flipping twice, the restriction of Theorem 1.2 to the elements of G not greater than g show that the preserved elements of \mathcal{U} force the recovery of the elements of \mathcal{D} less than g , and similarly that the preserved elements of \mathcal{D} force the recovery of the elements of \mathcal{U} greater than g . (And all other elements of \mathcal{D} and \mathcal{U} weren't affected by the flip). The second statement is immediate due to the order in which vertices are added to \mathcal{D} and \mathcal{U} . □

3.2. Independence posets

For G an acyclic directed graph, the *independence relations* on $\text{top}(G)$ are the reflexive and transitive closure of the relations $(\mathcal{D}, \mathcal{U}) < (\mathcal{D}', \mathcal{U}')$ if there is some $g \in \mathcal{U}$ such that $\text{flip}_g(\mathcal{D}, \mathcal{U}) = (\mathcal{D}', \mathcal{U}')$.

Lemma 3.3. *Independence relations are antisymmetric, and hence define an independence poset, denoted $\text{top}(G)$. Flips and cover relations of $\text{top}(G)$ coincide.*

Proof. Note that the relation is antisymmetric: any upwards flip from $(\mathcal{D}, \mathcal{U})$ at g introduces the new element g into \mathcal{D} , and the only way to remove it by a subsequent upwards flip is to flip at an element higher than g , which then introduces a new element not in \mathcal{D} . No such sequence of flips can therefore terminate at $(\mathcal{D}, \mathcal{U})$, and so the relation is antisymmetric.

By the definition of $\text{top}(G)$, every cover relation of G is induced by a flip. Suppose now that there is some $(\mathcal{D}, \mathcal{U}) \in \text{top}(G)$, and some g in \mathcal{U} , so that $(\mathcal{D}', \mathcal{U}') = \text{flip}_g(\mathcal{D}, \mathcal{U}) > (\mathcal{D}, \mathcal{U})$. What has to be checked is that there is no longer sequence of upward flips which also interpolates between $(\mathcal{D}, \mathcal{U})$ and $(\mathcal{D}', \mathcal{U}')$. Any upwards flip from $(\mathcal{D}, \mathcal{U})$ at h not below g would introduce the new element h into \mathcal{D} , which is not in \mathcal{D}' ; the only way to remove it by a subsequent upwards flip is to flip at an element higher than h , which then introduces a new element not in \mathcal{D}' into \mathcal{D} . No such sequence of flips can therefore terminate at $(\mathcal{D}', \mathcal{U}')$. Dually, any upwards flip at h not below g removes an element from \mathcal{U} which is contained in \mathcal{U}' , and the only way to restore it is to flip at a still lower element, which removes a different element of \mathcal{U}' from \mathcal{U} . Similarly, therefore, no such sequence of upwards flips can therefore terminate at $(\mathcal{D}', \mathcal{U}')$. \square

By Definition 1.4, the maximum element of $\text{top}(G)$ is the unique tight orthogonal pair $\hat{1}$ of the form (\mathcal{D}, \emptyset) , and its minimum element $\hat{0}$ is of the form (\emptyset, \mathcal{U}) . In particular, we see that $\text{top}(G)$ is connected. Figure 8 gives several examples of independence posets on various orientations of a path of length 3.

A *chain* in a poset is a sequence of elements $x_0 < x_1 < \dots < x_r$, of length r . The poset $\text{top}(G)$ has a maximal chain of length $|G|$ obtained by starting at $\hat{0}$ and flipping the elements of G in the order of a linear extension of G -order.

Lemma 3.4. *For G an acyclic directed graph with $g_1, \dots, g_{|G|}$ a linear extension of G -order, the sequence*

$$\hat{0} \triangleleft \text{flip}_{g_1}(\hat{0}) \triangleleft (\text{flip}_{g_2} \circ \text{flip}_{g_1})(\hat{0}) \triangleleft \dots \triangleleft (\text{flip}_{g_{|G|}} \circ \dots \circ \text{flip}_{g_1})(\hat{0}) = \hat{1}$$

is a maximal chain in $\text{top}(G)$.

Proof. Write $(\mathcal{D}_i, \mathcal{U}_i)$ for the i th element of the sequence. By Lemma 3.3, this sequence is unrefinable. By induction, after the i th step all elements of

\mathcal{U}_i lie above $\{g_1, \dots, g_i\}$ and all elements of \mathcal{D}_i are contained in $\{g_1, \dots, g_i\}$. Furthermore, since each \mathcal{U}_i is computed greedily in linear extension order, $g_{i+1} \in \mathcal{U}_i$. The sequence must end with $\hat{1}$, because the only way for all elements of $\mathcal{U}_{|G|}$ to lie above $\{g_1, \dots, g_{|G|}\}$ is for $\mathcal{U}_{|G|}$ to be empty. \square

Lemma 3.5. *Fix a linear extension ℓ of G -order. For any element $(\mathcal{D}, \mathcal{U}) \in \text{top}(G)$, there is a unique chain*

$$\hat{0} \triangleleft \text{flip}_{g_1}(\hat{0}) \triangleleft (\text{flip}_{g_2} \circ \text{flip}_{g_1})(\hat{0}) \triangleleft \dots \triangleleft (\text{flip}_{g_k} \circ \dots \circ \text{flip}_{g_1})(\hat{0}) = (\mathcal{D}, \mathcal{U})$$

such that $g_1 <_\ell g_2 <_\ell \dots <_\ell g_k$.

Proof. Starting with $(\mathcal{D}, \mathcal{U})$, flip elements h out of \mathcal{D} in reverse linear extension order. This does not add elements larger than h into \mathcal{D} . The process must therefore terminate with $\hat{0}$. Uniqueness follows from the fact that if we ever flip a lower element of \mathcal{D} than prescribed above, we will never be able to remove from \mathcal{D} the elements that we skipped over (without violating the constraint on the order of the flips). \square

Theorem 1.5. *Fix a linear extension ℓ of G -order. Flipping only in increasing order of ℓ gives a tree structure on the independent sets of G .*

Proof. Lemma 3.5 shows that by only permitting flips in the order of some fixed linear extension ℓ of G -order, we obtain a spanning tree of $\text{top}(G)$. \square

This tree structure is illustrated in Figure 6 for the Tamari lattice on 14 elements.

Proposition 3.6. *For G a directed acyclic graph, let G^* be the graph obtained by reversing all the edges in G . Then*

$$\begin{aligned} \text{top}(G) &\simeq \text{top}(G^*)^* \\ (\mathcal{D}, \mathcal{U}) &\mapsto (\mathcal{U}, \mathcal{D}), \end{aligned}$$

where $\text{top}(G^*)^*$ is the poset dual of $\text{top}(G^*)$.

Proof. Immediate. \square

3.3. Tight orthogonal pair recursion

For any $g \in G$, since $\{g\}$ is an independent set of G , by Theorem 1.2 there is a unique tight orthogonal pair m_g of the form $(\mathcal{D}, \{g\})$, and a unique tight orthogonal pair j_g of the form $(\{g\}, \mathcal{U})$. Write

$$(2) \quad \text{top}_g(G) := [\hat{0}, m_g] \text{ and } \text{top}^g(G) := [j_g, \hat{1}].$$

We say that $g \in G$ is *extremal* if it is a minimal or maximal element of G -order.

Lemma 3.7. *Let G be an acyclic directed graph. If g is an extremal element of G , then $\text{top}(G) = \text{top}_g(G) \sqcup \text{top}^g(G)$. Furthermore,*

- *If g is minimal, $(\mathcal{D}, \mathcal{U}) \in \text{top}_g(G)$ if and only if $g \in \mathcal{U}$, and*
- *If g is maximal, $(\mathcal{D}, \mathcal{U}) \in \text{top}^g(G)$ if and only if $g \in \mathcal{D}$.*

In particular, if $x \in \text{top}^g(G)$ and $y \in \text{top}_g(G)$, then $x \not\leq y$.

Proof. Suppose g is a minimal element of G and let $(\mathcal{D}, \mathcal{U}) \in \text{top}(G)$.

If $g \in \mathcal{U}$, then by successively flipping in any order all elements not equal to g that only cause us to move up in $\text{top}(G)$, we must end with the element m_g and so $(\mathcal{D}, \mathcal{U}) \in \text{top}_g(G)$. For certainly $g \in \mathcal{U}$ since it started out in \mathcal{U} and flipping at elements of G not equal to g does not remove g from \mathcal{U} , since g is minimal. And g is the only element of \mathcal{U} , or we would have flipped more elements out.

If $(\mathcal{D}, \mathcal{U}) \in \text{top}_g(G)$, then there is a sequence of flips that take us upwards to m_g . This sequence cannot remove g (since it could never be replaced), and so $g \in \mathcal{U}$.

If $g \in \mathcal{D}$, or if $g \notin \mathcal{D} \cup \mathcal{U}$, then we wish to successively flip in any order all elements not equal to g that only cause us to move down in $\text{top}(G)$. We claim that this must end with the element j_g so that $(\mathcal{D}, \mathcal{U}) \in \text{top}^g(G)$. Since the poset is finite, the process ends. Let $(\mathcal{D}', \mathcal{U}')$ be the end of this flip sequence. Then we claim $g \in \mathcal{D}'$ at this point. Note that \mathcal{D}' cannot contain any element $h \in G$ adjacent to g (or we would have flipped h out), so that g can be added to \mathcal{D}' —unless $g \in \mathcal{U}'$. But g cannot be in \mathcal{U}' , since it was not in \mathcal{U} , and only toggling at g could add it in.

The dual argument applies when g is maximal. □

Lemma 3.8. *For any element $g \in G$, if $g \in \mathcal{D}$ then $(\mathcal{D}, \mathcal{U}) \in \text{top}^g(G)$; and if $g \in \mathcal{U}$ then $(\mathcal{D}, \mathcal{U}) \in \text{top}_g(G)$.*

Proof. If $g \in \mathcal{D}$, then we construct a path downwards in $\text{top}(G)$ from $(\mathcal{D}, \mathcal{U})$ to j_g . First, flip out all elements greater than g in reverse linear extension order. This does not cause g to enter into \mathcal{U} (since all elements in \mathcal{U} less than h are fixed by flips). Since g can be added to \mathcal{D} at the end of this sequence of flips, it is indeed in \mathcal{D} (though it may not have stayed in it all the time during the sequence). Now flip out all the remaining elements of \mathcal{D} other than g . None of the remaining elements is above g , so this will neither remove g nor introduce elements in \mathcal{D} above g . The process necessarily terminates with j_g . Again, the statement for \mathcal{U} follows dually. □

Lemma 3.9. *Let G be a directed acyclic graph.*

- *If g is minimal and $(\mathcal{D}, \mathcal{U}) \in \text{top}_g(G)$, then $\text{flip}_g(\mathcal{D}, \mathcal{U}) = (\mathcal{D} \cup \{g\}, \mathcal{U}')$ for some \mathcal{U}' .*
- *If g is maximal and $(\mathcal{D}, \mathcal{U}) \in \text{top}^g(G)$, then $\text{flip}_g(\mathcal{D}, \mathcal{U}) = (\mathcal{D}', \mathcal{U} \cup \{g\})$ for some \mathcal{D}' .*

Proof. This follows from the definition of flip; when g is minimal, all of \mathcal{D} is preserved since every element of \mathcal{D} is not less than g . Similarly, when g is maximal, all of \mathcal{U} is preserved, since every element of \mathcal{U} is not greater than g . □

Theorem 3.10. *Let g be an extremal element of an acyclic directed graph G . Then*

$$\begin{aligned}
 (\mathcal{D}, \mathcal{U}) \mapsto (\mathcal{D}, \mathcal{U} \setminus \{g\}) \text{ is a bijection} & \begin{cases} \text{top}_g(G) \simeq \text{top}(G_g^\circ) & \text{if } g \text{ minimal} \\ \text{top}_g(G) \simeq \text{top}(G_g) & \text{if } g \text{ maximal} \end{cases} \\
 (\mathcal{D}, \mathcal{U}) \mapsto (\mathcal{D} \setminus \{g\}, \mathcal{U}) \text{ is a bijection} & \begin{cases} \text{top}^g(G) \simeq \text{top}(G_g) & \text{if } g \text{ minimal} \\ \text{top}^g(G) \simeq \text{top}(G_g^\circ) & \text{if } g \text{ maximal} \end{cases}
 \end{aligned}$$

Proof. We only prove the results for g minimal, the case for g maximal being analogous. We first show $\text{top}_g(G) \simeq \text{top}(G_g^\circ)$; by Lemma 3.7, $(\mathcal{D}, \mathcal{U}) \in \text{top}_g(G)$ if and only if $g \in \mathcal{U}$. But since $(\mathcal{D}, \mathcal{U})$ is a tight orthogonal pair, no element of \mathcal{D} or of \mathcal{U} can be adjacent to g , from which we conclude the result by definition of G_g° . We now show $\text{top}^g(G) \simeq \text{top}(G_g)$; since $\text{top}(G) = \text{top}_g(G) \sqcup \text{top}^g(G)$, by Lemma 3.7 elements of $\text{top}^g(G)$ consist of those tight orthogonal pairs of G with either $g \in \mathcal{D}$ or $g \notin \mathcal{U} \cup \mathcal{D}$. Each tight orthogonal pair of G_g can be uniquely extended to such a tight orthogonal pair. □

4. Trim lattices and maximal orthogonal pairs

4.1. Extremal lattices

An *extremal lattice* is a lattice whose longest chain is of length equal to the number of its join irreducible elements and to the number of its meet irreducible elements. As motivation for our main result of this section, we have the following easy statement (we will refine it in Theorem 1.6).

Lemma 4.1. *If $\text{top}(G)$ is a lattice, then it is an extremal lattice.*

Proof. Note that a lattice with a chain of length n (i.e. with $n + 1$ elements) must have at least n join-irreducible elements and at least n meet-irreducible elements (since each element of the chain is the join of the join-irreducibles beneath it and the meet of the meet-irreducibles above it). Suppose $\text{top}(G)$ is a lattice; since it only has $|G|$ join and $|G|$ meet irreducible elements, and since it has a chain of length $|G|$ by Lemma 3.4, it is extremal. \square

In [TW17], we represented extremal lattices in the following way, following a construction of Markowsky [Mar92]. Any acyclic directed graph G gives rise to an extremal lattice $\mathcal{L}(G)$, as follows: for $X, Y \subseteq G$ with $X \cap Y = \emptyset$, we say (X, Y) is an *orthogonal pair* if there is no edge from any $i \in X$ to any $k \in Y$, and we say it is a *maximal orthogonal pair* if X and Y are maximal with that property. Clearly, to each $Y \subseteq G$, there is at most one X such that (X, Y) is a maximal orthogonal pair (and dually). Then the extremal lattice $\mathcal{L}(G)$ is equivalently given by *either* of

$$\begin{aligned} (X, Y) \leq (X', Y') &\text{ if and only if } X \subseteq X', \text{ or} \\ (X, Y) \leq (X', Y') &\text{ if and only if } Y' \subseteq Y. \end{aligned}$$

Furthermore, the join is computed by intersecting the second terms, while meet is given by the intersection of the first terms. If x is an element of an extremal lattice $\mathcal{L}(G)$ with corresponding maximal orthogonal pair (X, Y) , we write $x_{\mathcal{J}} = X$ and $x_{\mathcal{M}} = Y$ —that is, $x_{\mathcal{J}}$ corresponds to the join-irreducible elements below x , while $x_{\mathcal{M}}$ corresponds to the meet-irreducible elements above x . Two examples are given in Figure 5.

Conversely, we can associate an acyclic directed graph $G(\mathcal{L})$ to any extremal lattice called its *Galois graph* with the property that $\mathcal{L}(G(\mathcal{L})) \simeq \mathcal{L}$. We refer to [TW17] for further details, including Markowsky’s generalization of Birkhoff’s fundamental theorem of distributive lattices to extremal lattices.

4.2. Trim lattices

An element x of a lattice \mathcal{L} is called *left modular* if for any $y \leq z$ we have the equality

$$(y \vee x) \wedge z = y \vee (x \wedge z).$$

A lattice is called *left modular* if it has a maximal chain of left modular elements.

A *trim lattice* is an extremal left-modular lattice. We have already shown that if an independence poset is a lattice, then it is extremal. Our goal is to prove that it is actually trim.

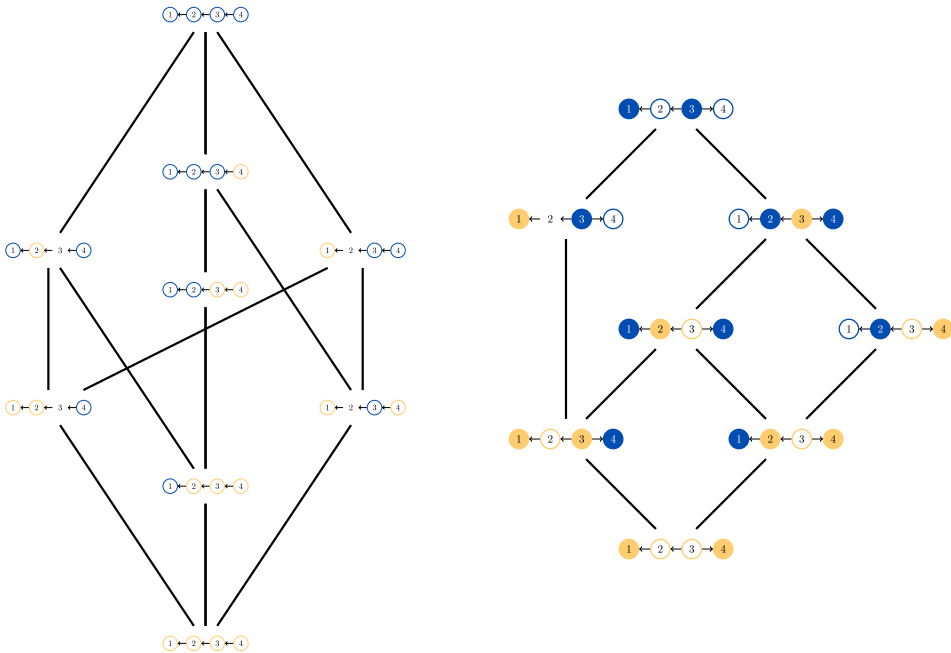


Figure 5: On the left is the extremal lattice of maximal orthogonal sets of a directed path of length 4, ordered by inclusion of first component; the sets x_f are indicated by a blue border, while the sets x_m are drawn with a yellow border. Because it is not trim, this extremal lattice is not isomorphic to the top left example in Figure 3 (which is not even a lattice). On the right is the (trim) extremal lattice for a second orientation, which does coincide with the top right example in Figure 3 by Theorem 4.9; maximal orthogonal pairs are indicated by the color of the border, while tight orthogonal pairs are indicated by the color of the filling.

We say that a relation $y < z$ in an extremal lattice $\mathcal{L}(G)$ is *overlapping* if

$$y_m \cap z_f \neq \emptyset.$$

Theorem 4.2 ([TW17, Theorem 3.4]). *An extremal lattice $\mathcal{L}(G)$ is trim if and only if every relation is overlapping if and only if every cover relation is overlapping.*

If a cover relation is overlapping, then it overlaps in a unique element. We may define the *downward* and *upward labels* of $y \in \mathcal{L}(G)$ as

$$D(y) := \{\text{the unique element of } x_m \cap y_f : \text{all } x \text{ such that } x < y\} \text{ and}$$

$U(y) := \{\text{the unique element of } y_{\mathcal{M}} \cap z_{\mathcal{J}} : \text{all } z \text{ such that } y \prec z\}.$

In a trim lattice $\mathcal{L}(G)$, there is a unique meet-irreducible element m_g with $U(m_g) = \{g\}$, and a unique join-irreducible element j_g with $D(j_g) = \{g\}$. We proved the following recursive properties of trim lattices in [TW17], which are exactly analogous to Lemma 3.7 and Theorem 3.10 and will allow us to relate maximal orthogonal pairs and tight orthogonal pairs.

Theorem 4.3 ([TW17, Lemma 3.10, Proposition 3.11, Proposition 3.12]). *Let g be minimal in an acyclic directed graph G , and write $\mathcal{L}_g(G) := [\hat{0}, m_g]$ and $\mathcal{L}^g(G) := [j_g, \hat{1}]$. Then*

1. $\mathcal{L}(G) = \mathcal{L}_g(G) \sqcup \mathcal{L}^g(G)$,
2. $\mathcal{L}^g(G) \simeq \mathcal{L}(G_g)$,
3. $\mathcal{L}_g(G) \simeq \mathcal{L}(G_g^\circ)$, and
4. an element $x \in \mathcal{L}_g(G)$ if and only if $g \in U(x)$.

We recall that the downward and upward labels actually associate two independent sets to each element of \mathcal{L} .

Theorem 4.4 ([TW17, Corollary 5.6]). *For \mathcal{L} a trim lattice, D and U are both bijections from \mathcal{L} to the set of independent sets of $G(\mathcal{L})$.*

We will improve this in Section 4.4—taking both the downward and upward labels together give a tight orthogonal pair.

4.3. Trim lattices to independence posets

The next subsections relate trim lattices and independence posets, simultaneously generalizing the bijections between order ideals and antichains, and between Coxeter-sortable elements in a finite Coxeter group and the corresponding noncrossing partitions.

We show that $\text{top}(G) \simeq \mathcal{L}(G)$ under certain conditions on $\text{top}(G)$ or on $\mathcal{L}(G)$. If $\mathcal{L}(G)$ is a lattice, then $\text{top}(G) \simeq \mathcal{L}(G)$ (Theorem 1.7). Similarly, if $\text{top}(G)$ is a lattice, then also $\text{top}(G) \simeq \mathcal{L}(G)$ (Corollary 4.7). We also show that if $\text{top}(G)$ is a lattice, then it is a trim lattice (Theorem 1.6).

Theorem 1.7. *If $\mathcal{L}(G)$ is trim, then $\mathcal{L}(G) \simeq \text{top}(G)$.*

Proof. We argue by induction on $|G|$. Let g be minimal in G . By Theorem 4.3, since $\mathcal{L}(G)$ is trim, $\mathcal{L}(G) = \mathcal{L}_g(G) \sqcup \mathcal{L}^g(G)$. Similarly, $\text{top}(G) = \text{top}_g(G) \sqcup \text{top}^g(G)$ by Theorem 3.10. By induction, we conclude the isomorphism on each of G_g° and G_g : $\mathcal{L}_g(G) \simeq \mathcal{L}(G_g^\circ) \simeq \text{top}(G_g^\circ) \simeq \text{top}_g(G)$ and $\mathcal{L}^g(G) \simeq \mathcal{L}(G_g) \simeq \text{top}(G_g) \simeq \text{top}^g(G)$. Moreover, this induction respects the labelling of cover relations by the overlapping element (for the mops) and the element flipped (for the tops).

We have only to show now that the cover relations between $\mathcal{L}_g(G)$ and $\mathcal{L}^g(G)$ are the same as those for $\text{top}_g(G)$ and $\text{top}^g(G)$. But each element in $\mathcal{L}_g(G)$ and $\text{top}_g(G)$ has an edge up by Theorem 4.3 (4) and Lemma 3.9 (labelled by g). The element $x \in \mathcal{L}_g(G)$ is paired with the unique element $x' \in \mathcal{L}^g(G)$ satisfying $D(x') = D(x) \cup \{g\}$ by [TW17, Lemma 3.15]. It is evident from the definition of flip that the same rule describes how to pair elements in $\text{top}_g(G)$ with elements of $\text{top}^g(G)$. \square

We now analyze what can be deduced from the fact that $\text{top}(G)$ is a lattice. To begin with, we show that, if it is a lattice, it is necessarily trim.

Theorem 1.6. *If $\text{top}(G)$ is a lattice then it is a trim lattice.*

Proof. Suppose that $\text{top}(G)$ is a lattice. By Lemma 4.1, it is therefore an extremal lattice. We wish to show that every cover relation is overlapping, so that $\text{top}(G)$ is trim by Theorem 4.2. Consider the representation of $\text{top}(G)$ as maximal orthogonal sets, which gives the correspondence $(\mathcal{D}, \mathcal{U})$ with (X, Y) , where X is the set of join-irreducible elements below $(\mathcal{D}, \mathcal{U})$ and Y is the set of meet-irreducible elements above $(\mathcal{D}, \mathcal{U})$. By Lemma 3.8, we therefore have that $\mathcal{D} \subseteq X$ and $\mathcal{U} \subseteq Y$. But if $g \in \mathcal{U}$ and $\text{flip}_g(\mathcal{D}, \mathcal{U}) = (\mathcal{D}', \mathcal{U}')$ is a cover with corresponding maximal orthogonal sets $(X, Y) \triangleleft (X', Y')$, then $\mathcal{D}' \cap \mathcal{U}' = \{g\} \subseteq X' \cap Y'$ and hence the cover relation is overlapping. \square

The following lemma describes a situation in which it is possible to start from an orthogonal pair of independent sets and produce a tight orthogonal pair by adding elements to the two sets. It is needed for the proof of the next theorem.

Lemma 4.5. *Let $(\mathcal{D}', \mathcal{U}')$ be an orthogonal pair of independent sets, and suppose that no element of \mathcal{D}' is below an element of \mathcal{U}' in G -order. Then there exists a tight orthogonal pair of independent sets $(\mathcal{D}, \mathcal{U})$ such that $\mathcal{D} \supseteq \mathcal{D}'$ and $\mathcal{U} \supseteq \mathcal{U}'$.*

Proof. In a linear extension ℓ of G -order, greedily add any element to \mathcal{U}' which can be added, subject to the condition that the set remain orthogonal to \mathcal{D}' . Let \mathcal{U} be the resulting set. Define \mathcal{D} to be the independent set given by Algorithm 2 such that $(\mathcal{D}, \mathcal{U})$ is tight orthogonal.

All that is necessary to show is that \mathcal{D} contains \mathcal{D}' . Suppose that, as we are constructing \mathcal{D} using Algorithm 2, that there is some element x of \mathcal{D}' which we do not add. The reason we do not add it must be because there is some $y \rightarrow x$ such that we do add y to \mathcal{D} . This element y necessarily has no edge to any element of \mathcal{U} . Since, by the hypothesis on the relative order

of \mathcal{D}' and \mathcal{U}' , there is also no edge from any element of \mathcal{U}' to y , we would in fact have added y into \mathcal{U} , which is a contradiction. Therefore \mathcal{D} contains \mathcal{D}' as desired. \square

Note that without the assumption on the relative order of \mathcal{D}' and \mathcal{U}' , the conclusion of this lemma is false, as demonstrated by the case of the linearly oriented path $(1) \leftarrow (2) \leftarrow (3) \leftarrow (4)$ on four vertices: there is no tight orthogonal pair with $\mathcal{D} \supseteq \{2\}$ and $\mathcal{U} \supseteq \{4\}$ (see the top left of Figure 3).

Assuming that $\text{top}(G)$ is a lattice, and thus trim by Theorem 1.6, it is of the form $\mathcal{L}(H)$ for some H . We show that in this case $G \simeq H$.

Theorem 4.6. *If $\text{top}(G) \simeq \mathcal{L}(H)$, then $G \simeq H$.*

Proof. The join-irreducibles and meet-irreducibles of the extremal lattice $\mathcal{L}(H)$ are canonically identified, by [TW17, Proposition 2.7], and are also identified with the vertices of H . The vertices of G are likewise canonically identified with the join-irreducibles and the meet-irreducibles of $\text{top}(G)$; and the identification of the join-irreducibles and meet-irreducibles defined in this way is the same as the identification coming from $\mathcal{L}(H)$. The vertices of G and of H are therefore naturally identified.

Recall from [TW17] that the *spine* of an extremal lattice $\mathcal{L}(H)$ is the collection of elements lying on maximal-length chains; the spine is a distributive sublattice, isomorphic to the distributive lattice $J(P)$ where P is the poset corresponding to H -order. In $\text{top}(G)$ the maximal chains of the spine correspond to sequences of flips in which every vertex is flipped once. By Lemma 3.4 this can certainly be done in any linear extension of G -order; in principle, other orders could also be possible. This shows that G -order is at least as strong as H -order.

If $g_1 \rightarrow g_2$ is an edge of H , then g_1 is above g_2 in H -order, so g_1 is above g_2 in G -order, and in particular, g_1 is not below g_2 in G -order. Suppose, seeking a contradiction, that there were no edge $g_1 \rightarrow g_2$ in G . In this case $(\{g_1\}, \{g_2\})$ is an orthogonal pair, and, by Lemma 4.5, there is a top $(\mathcal{D}, \mathcal{U})$ with $g_1 \in \mathcal{D}$ and $g_2 \in \mathcal{U}$, so $j_{g_1} \leq m_{g_2}$, contradicting the fact that there is an edge $g_1 \rightarrow g_2$ in H .

It follows that the edges of H are a subset of the edges of G . Since $\mathcal{L}(H)$ is trim, by Theorem 1.7, $\mathcal{L}(H) \simeq \text{top}(H)$. If the edges of H were a strict subset of the edges of G , G would have more independent subsets than H , so $\text{top}(G)$ would have more vertices than $\mathcal{L}(H)$, which is a contradiction. Thus $G \simeq H$. \square

Theorem 1.6 and Theorem 4.6 together imply an isomorphism of tops and mops when $\text{top}(G)$ is a lattice.

Corollary 4.7. *If $\text{top}(G)$ is a lattice, then $\text{top}(G) \simeq \mathcal{L}(G)$.*

Proof. Since $\text{top}(G)$ is a lattice, it is a trim lattice, and hence is $\mathcal{L}(H)$ for some directed acyclic graph H . But now $H \simeq G$ by Theorem 4.6. \square

4.4. Bijections

Theorem 4.8. *Let $\mathcal{L}(G)$ be a trim lattice and $x \in \mathcal{L}(G)$. Then $\phi(x) = (\text{D}(x), \text{U}(x))$ is a tight orthogonal pair.*

Proof. For $g \in G$, write $j_g \in \mathcal{L}(G)$ for the unique join irreducible with $\text{D}(j_g) = \{g\}$ and $m_g \in \mathcal{L}(G)$ for the unique meet irreducible element with $\text{U}(m_g) = \{g\}$. Given a collection of join-irreducible elements $J = \{j_{g_1}, \dots, j_{g_r}\}$ and meet-irreducible elements $M = \{m_{g'_1}, \dots, m_{g'_s}\}$, then there is no edge from an element of the set $\{g_1, \dots, g_r\}$ to an element of the set $\{g'_1, \dots, g'_s\}$ if and only if every element of J is below every element of M if and only if the join of J is below the meet of M .

By Theorem 4.4, every $x \in \mathcal{L}(G)$ is associated to a pair of independent sets $(\text{D}(x), \text{U}(x))$. By [TW17, Proposition 4.1],

$$\bigvee_{g \in \text{D}(x)} j_g = x = \bigwedge_{g \in \text{U}(x)} m_g.$$

By the previous paragraph, there are no edges from any element of $\text{D}(x)$ to any element of $\text{U}(x)$, so the pair $(\text{D}(x), \text{U}(x))$ is orthogonal.

We now argue that $(\text{D}(x), \text{U}(x))$ is tight. If we increased any element of $\text{D}(x)$ while staying independent, this new set would correspond to $\text{D}(x')$ of some $x' \in \mathcal{L}(G)$ (by Theorem 4.4). This x' would be strictly greater than x —in particular, it would no longer be below every m_g for $g \in \text{U}(x)$, and so $(\text{D}(x), \text{U}(x))$ would no longer be orthogonal. \square

Theorem 4.9. *If $\mathcal{L}(G)$ is trim—or, equivalently, if $\text{top}(G)$ is a lattice—then ϕ is an isomorphism $\mathcal{L}(G) \simeq \text{top}(G)$.*

Proof. We show ϕ is an isomorphism by induction on $|G|$: to this end, let g be minimal in G . Recall that $\mathcal{L}(G) = \mathcal{L}_g(G) \sqcup \mathcal{L}^g(G)$ and $\text{top}(G) = \text{top}_g(G) \sqcup \text{top}^g(G)$. Since $\mathcal{L}_g(G) \simeq \mathcal{L}(G_g^\circ)$ and $\text{top}_g(G) \simeq \text{top}(G_g^\circ)$ and $x \in \mathcal{L}_g(G)$ if and only if $g \in \text{U}(x)$ and $(\mathcal{D}, \mathcal{U}) \in \text{top}_g(G)$ if and only if $g \in \mathcal{U}$, we conclude that ϕ is an isomorphism from $\mathcal{L}_g(G) \simeq \text{top}_g(G)$. Similarly, we conclude that ϕ is an isomorphism from $\mathcal{L}^g(G) \simeq \text{top}^g(G)$, since $\mathcal{L}^g(G) \simeq \mathcal{L}(G_g)$ and $\text{top}^g(G) = \text{top}(G_g)$.

We have only to show now that the cover relations between $\mathcal{L}_g(G)$ and $\mathcal{L}^g(G)$ are the same as those for $\text{top}_g(G)$ and $\text{top}^g(G)$. But this now follows from the same argument as in Theorem 1.7. \square

The recurrence used in Theorem 4.9 can be used to show that ϕ has the following alternative description. Given $(X, Y) \in \mathcal{L}(G)$, construct $(\mathcal{D}, \mathcal{U}) \in \text{top}(G)$ as follows: \mathcal{D} is obtained from X by greedily choosing elements from X in a reverse linear extension order of G -order so that \mathcal{D} remains an independent set; \mathcal{U} is obtained from Y by greedily choosing elements from Y in a linear extension order so that \mathcal{U} remains an independent set.

The inverse of the map ϕ is constructed in the following theorem.

Theorem 4.10. *Given $(\mathcal{D}, \mathcal{U}) \in \text{top}(G)$, set $X = \mathcal{D}$, $Y = \mathcal{U}$, and define*

θ_1 : add to Y all elements of $G \setminus X$ with no arrow from an element of X .

θ_2 : add to X all elements of $G \setminus Y$ with no arrow to an element of Y .

Then both $\theta_1 \circ \theta_2$ and $\theta_2 \circ \theta_1$ are injections $\text{top}(G) \hookrightarrow \mathcal{L}(G)$. Moreover, if $\text{top}(G)$ is a lattice, then they coincide and are inverses to ϕ .

Proof. It is clear that both $\theta_1 \circ \theta_2$ and $\theta_2 \circ \theta_1$ have images in $\mathcal{L}(G)$: $\theta_2 \circ \theta_1$, we cannot add anything to X , but we also cannot add anything to Y because it was chosen to be maximal with respect to a subset of X —and similarly for $\theta_1 \circ \theta_2$.

For $\theta_2 \circ \theta_1$, we first reconstruct \mathcal{D} from X by greedily choosing elements in X in a reverse linear extension order— Y contains those elements that don't have an arrow from an element of \mathcal{D} , so as we scan through X , if we see an element without an arrow from our reconstruction of \mathcal{D} , we must add it to \mathcal{D} since otherwise it would be in Y . The set \mathcal{U} is now determined by Theorem 1.2. A similar argument works for $\theta_1 \circ \theta_2$ —we first reconstruct \mathcal{U} from Y by greedily choosing elements in Y in a linear extension order, and then \mathcal{D} is determined.

Finally, when $\text{top}(G)$ is a lattice, these coincide with the alternative description of ϕ given after Theorem 4.9. □

In particular, the number of tight orthogonal pairs is always less than or equal to the number of maximal orthogonal pairs. Theorem 4.9 is illustrated in Figure 5.

4.5. Conditions on independence lattices

Overlapping relations in a trim lattice allow us to give a graph-theoretic condition on G for when $\text{top}(G)$ is a lattice.

Theorem 4.11. *For G an acyclic directed graph, $\text{top}(G)$ is a lattice if and only if G has no partition $G = X_1 \sqcup X_2 \sqcup X_3 \sqcup X_4 \sqcup Z$ such that*

- (i) all the X_i are non-empty (but Z may be empty)*
- (ii) every element of X_3 has an edge from X_4 and to X_2*

- (iii) every element of X_2 has an edge to X_1 and from X_3
- (iv) there are no edges from X_4 to X_2 , from X_4 to X_1 , or from X_3 to X_1
- (v) every element of Z has an edge from X_4 and an edge to X_1 .

Proof. Suppose $\mathcal{L}(G)$ is not trim, so that we have a cover relation $(X, Y) \triangleleft (X', Y')$ with $Y \cap X' = \emptyset$. Let $X_4 = X$, $X_3 = X' \setminus X$, $X_2 = Y \setminus Y'$, $X_1 = Y'$, and let Z contain the remaining elements of G . Every element of X_3 has an edge from some element of X_4 (since otherwise Y should be bigger) and an edge to some element of X_2 (since otherwise X should be bigger). Similarly, every element of X_2 has an edge to some element of X_1 and an edge from some element of X_3 . There are no edges from X_4 to X_2 or to X_1 (this would contradict orthogonality of (X, Y)) or from X_3 to X_1 (this would contradict orthogonality of (X', Y')). Every element of Z , however, has a edge from X_4 and a edge to X_1 .

Conversely, suppose $\mathcal{L}(G)$ is trim and consider any partition of G into five sets X_4, X_3, X_2, X_1, Z satisfying conditions (i), (ii), (iii) and (iv). We will show that condition (v) is violated. Remove all edges from X_3 to X_2 , so that $(X_4 \cup X_3, X_2 \cup X_1)$ is an orthogonal pair, and can be extended to a maximal orthogonal pair (X, Y) . Adding the edges from X_3 to X_2 back in, we may find maximal orthogonal pairs of the form $(X \setminus X_3, Y')$ and $(X', Y \setminus X_2)$. These pairs are comparable, and their overlap $X' \cap Y'$ consists only of vertices of G with an edge from X_3 , an edge to X_2 , and no edge from X_4 or to X_1 . \square

4.6. Independence posets that are not lattices

In this section we highlight several differences between the behavior of independence posets that are (trim) lattices, and independence posets that are not lattices.

Independence posets break the rigidity of the spine (the elements lying on maximal-length chains) of a trim lattice: although linear extensions of G -order still index *certain* maximal chains of $\text{top}(G)$, there can now be other maximal chains. In fact, the length of a maximal length chain in $\text{top}(G)$ can be strictly greater than $|G|$ —that is, the same element can be flipped from \mathcal{U} to \mathcal{D} multiple times in the same chain.

Example 4.12. We leave it to the reader to check that the length of the longest chain in independence poset associated to a directed path on five vertices is not five, but six.

Furthermore, although an independent poset that is a lattice arises from a unique Galois graph—since an extremal lattice determines its Galois




Figure 6: The Tamari lattice with 14 elements, realized as an independence poset. The thick blue edges indicate the tree structure provided by the natural labelling, giving an efficient method to generate all independent sets of the underlying graph. The filling of the vertices of the graph specify the tight orthogonal pairs, while the color of the boundaries specify the maximal orthogonal pair.




Figure 7: An independence poset $\text{top}(G)$ with the property that $\text{top}(G) \simeq \text{top}(G)^*$ but $G \not\simeq G^*$.

graph, up to isomorphism—Proposition 4.13 shows that uniqueness does not hold for general independence posets.

Proposition 4.13. *Nonisomorphic directed acyclic graphs can give isomorphic independence posets.*

Proof. Let G be as in Figure 7—observe that $\text{top}(G) \simeq \text{top}(G)^*$, but that $G^* \not\simeq G$. By Proposition 3.6, this gives the desired example. \square

As a referee observed, although trim lattices are always EL-shellable with Möbius function taking only the values $-1, 0,$ or 1 , neither of these proper-

ties are true for independence posets. For example, the Möbius function on the independence poset built from a linearly ordered path with six vertices attains a maximum value of 4; one can check that the independence poset built from a linearly order path with four vertices is already not shellable.

5. Toggles

Fix an undirected graph G and any element $g \in G$. In [Str16, Section 3.6] and in [JR17], a *toggle* of an independent set \mathcal{F} of graph G is defined by Equation (3)

$$(3) \quad \text{tog}_g(\mathcal{F}) = \begin{cases} \mathcal{F} \cup \{g\} & \text{if } g \notin \mathcal{F} \text{ and } \mathcal{F} \cup \{g\} \text{ is an independent set,} \\ \mathcal{F} \setminus \{g\} & \text{if } g \in \mathcal{F}, \\ \mathcal{F} & \text{otherwise.} \end{cases}$$

Toggles appear naturally in independence posets through an operation similar to quiver mutation on the underlying directed graph G , but are defined in our context only when g is an extremal element of G . For g an extremal element of G , the *toggle* of the graph G at g is the acyclic directed graph $\text{tog}_g(G)$ obtained by reversing all edges incident to g .

We have chosen the term “toggle” for consistency with [SW12, Str16]. Although [Str16, JR17] define toggles as bijections on a fixed set of independent sets, our operation comes from changing the orientation of the underlying directed graph G —and hence the underlying set of tops is not fixed. We show below that we can restrict to one component of the tops to recover their bijections on independent sets.

It is clear that $\text{tog}_g^2(G) = G$. When g is extremal, by Lemma 3.7 no element of $\text{top}^g(G)$ lies below an element of $\text{top}_g(G)$ in $\text{top}(G)$. The effect of toggling at g is to reverse the roles of $\text{top}_g(G)$ and $\text{top}^g(G)$ —roughly, transporting $\text{top}_g(G)$ above $\text{top}^g(G)$. This relationship between $\text{top}(G)$ and $\text{top}(\text{tog}_g(G))$ is summarized in Theorem 5.1, and is illustrated in Figure 8.

Theorem 5.1. *Let g be a minimal element of an acyclic directed graph G . Then*

- $(\mathcal{D}, \mathcal{U}) \mapsto (\mathcal{D} \cup \{g\}, \mathcal{U} \setminus \{g\})$ is a bijection $\text{top}_g(G) \simeq \text{top}^g(\text{tog}_g(G))$,
- $(\mathcal{D}, \mathcal{U}) \mapsto (\mathcal{D} \setminus \{g\}, \mathcal{U}')$ is a bijection $\text{top}^g(G) \simeq \text{top}_g(\text{tog}_g(G))$.

Let g be a maximal element of an acyclic directed graph G . Then

- $(\mathcal{D}, \mathcal{U}) \mapsto (\mathcal{D}', \mathcal{U} \setminus \{g\})$ is a bijection $\text{top}_g(G) \simeq \text{top}^g(\text{tog}_g(G))$,
- $(\mathcal{D}, \mathcal{U}) \mapsto (\mathcal{D} \setminus \{g\}, \mathcal{U} \cup \{g\})$ is a bijection $\text{top}^g(G) \simeq \text{top}_g(\text{tog}_g(G))$.




Figure 8: Each poset is an independence poset on an orientation of a path of length 3. Taken together, these posets represent a sequence of toggles. Each toggle induces a bijection between $\text{top}(G)$ and $\text{top}(\text{tog}_g(G))$, which we keep track of by assigning letters to the elements of the posets. Toggling each element of the directed graph underlying the independence poset at the top left (in reverse linear extension order) recovers the same independence poset on the bottom left, permuting its elements as $(a, b, c, d, e) \xrightarrow{\text{tog}_{321}} (e, d, b, c, a)$. This coincides with rowmotion.

Proof. This follows immediately from Theorem 3.10. □

For $g \in G$ extremal, we also write tog_g for the bijection of Theorem 5.1 from $\text{top}(G) \rightarrow \text{top}(\text{tog}_g(G))$, and call it a *toggle*. By Theorem 5.1, we have that $\text{tog}_g^2(\mathcal{D}, \mathcal{U}) = (\mathcal{D}, \mathcal{U})$.

Thus, if we limit ourselves to toggling only at minimal elements of G and keeping track of the first components \mathcal{D} , then Equation (3) allows us

to compute tog_g on independent sets. Equation (3) also applies if we choose only maximal elements of G and look at only the second components \mathcal{U} in Theorem 5.1.

6. Rowmotion on independence posets

By Theorem 1.2, any independent set \mathcal{F} can be completed to a tight orthogonal pair in exactly two ways—the first as $(\mathcal{F}, \mathcal{U})$ and the second as $(\mathcal{D}, \mathcal{F})$. *Rowmotion* sends the first of these to the second (see Equation (1)). The purpose of this section is to give two additional ways to compute it: one using flips, and one using toggles.

Definition 6.1. For G an acyclic directed graph, we say that rowmotion on $\text{top}(G)$ can be computed

- *in slow motion* if $\text{row} = \prod_{g \in \ell} \text{flip}_g$, and
- *by deformation* if $\text{row} = \prod_{g \in \ell'} \text{tog}_g$,

where ℓ is any linear extension of G -order and ℓ' is any reverse linear extension of G -order.

Theorem 1.8(i). *For G an acyclic directed graph, rowmotion on $\text{top}(G)$ can be computed in slow motion.*

Proof. We follow the same proof as [TW17, Theorem 1.1]. Let g be minimal in G .

Case I: $(\mathcal{D}, \mathcal{U}) \in \text{top}_g(G)$. At the first step when calculating $\prod_{h \in \ell} \text{flip}_h$, we walk to $(\mathcal{D}', \mathcal{U}') = \text{flip}_g(\mathcal{D}, \mathcal{U})$. So $\mathcal{D}' = \mathcal{D} \cup \{g\}$ by Lemma 3.9 and $(\mathcal{D}', \mathcal{U}') \in \text{top}^g(G)$ by Lemma 3.7. Applying the rest of the flips to $(\mathcal{D}', \mathcal{U}')$ has the effect of applying $\prod_{\substack{h \in \ell \\ h \neq g}} \text{flip}_h$ in $\text{top}(G_g)$. (Note that the only flips in $\text{top}(G)$ leaving $\text{top}^g(G)$ are those of the form flip_g , so they will never be taken.) By induction, we obtain an element $(\mathcal{D}'', \mathcal{U}'') \in \text{top}(G_g)$ such that $\mathcal{U}'' = \mathcal{D}' \setminus \{g\}$. Passing back to $\text{top}^g(G)$ does not change \mathcal{U}'' , and we conclude the result.

Case II: $(\mathcal{D}, \mathcal{U}) \in \text{top}^g(G)$ and $g \in \mathcal{D}$. At the first step when calculating $\prod_{h \in \ell} \text{flip}_h$, we walk to $(\mathcal{D}', \mathcal{U}') = \text{flip}_g(\mathcal{D}, \mathcal{U})$. Since $g \in \mathcal{U}'$, $(\mathcal{D}', \mathcal{U}') \in \text{top}_g(G)$ by Lemma 3.7 and $\mathcal{D}' \cup \{g\} = \mathcal{D}$ by Lemma 3.9. Applying the rest of the flips to $(\mathcal{D}', \mathcal{U}')$ has the effect of applying $\prod_{\substack{h \in \ell \\ h \neq g}} \text{flip}_h$ in $\text{top}_g(G)$. (Note that the only flips in $\text{top}(G)$ leaving $\text{top}_g(G)$ are those of the form flip_g , so they will never be taken.) By induction, we obtain an element $(\mathcal{D}'', \mathcal{U}'') \in \text{top}_g(G)$ such that $\mathcal{U}'' = \mathcal{D}' \cup \{g\}$. But $\mathcal{D}' \cup \{g\} = \mathcal{D}$.

Case III: $(\mathcal{D}, \mathcal{U}) \in \text{top}^g(G)$ and $g \notin \mathcal{D}$. By assumption, flipping at g has no effect. As in Case I, the rest of the flips have the effect of applying $\prod_{\substack{h \in \ell \\ h \neq g}} \text{flip}_h$ in $\text{top}^g(G)$. By induction, we obtain an element $(\mathcal{D}', \mathcal{U}') \in \text{top}^g(G)$ such that $\mathcal{U}' = \mathcal{D}$. \square

Theorem 1.8(ii). *For G an acyclic directed graph, rowmotion on $\text{top}(G)$ can be computed by deformation.*

Proof. Since every element of G is toggled, every edge is flipped twice, and so $\prod_{g \in \ell'} \text{tog}_g(G) = G$. By definition, tog_g takes all elements with $g \in \mathcal{D}$ and converts them to elements with $g \in \mathcal{U}$. Since we are toggling in $\ell' = g_1, \dots, g_{|G|}$ order, at any step we have moved all $\{g_1, \dots, g_i\} \cap \mathcal{D}$ from \mathcal{D} to \mathcal{U} . Although this introduces some new elements into \mathcal{D} , these only involve $\{g_1, \dots, g_i\}$ and so are never moved to \mathcal{U} in a subsequent step. Thus, each element has its *original* set \mathcal{D} converted to its new \mathcal{U} , which is the definition of rowmotion given in Equation (1). \square

7. Representation theory

In this section, we show that the combinatorics of independence lattices arises naturally in representation theory. In short, the torsion/torsion-free pairs of certain acyclic finite-dimensional algebras correspond to maximal orthogonal pairs, while their 2-term simpleminded collections recover tight orthogonal pairs. The setting in which we work, while special from the point of view of representation theory, includes many interesting examples, such as all quotients of Dynkin path algebras.

7.1. Representation-finite directed algebras

Let k be a field, and A a finite-dimensional k -algebra. Suppose further that the module category $\text{mod } A$ is directed, i.e., there is no sequence of pairwise non-isomorphic indecomposable modules M_1, \dots, M_r for $r > 1$ with non-zero morphisms from M_i to M_{i+1} and from M_r to M_1 . Under this hypothesis, all indecomposable modules have some properties which usually only belong to a subset of indecomposables.

A module is called a *brick* if its endomorphism ring is a division algebra. Note that a brick is necessarily indecomposable. A module M is called τ -*rigid* if

$$\text{Hom}(M, \tau M) = 0,$$

where τ is the Auslander-Reiten translation.

Proposition 7.1. *Let A be a finite-dimensional k -algebra such that $\text{mod } A$ has no cycles. Then all indecomposable A -modules are τ -rigid bricks.*

Proof. Suppose that M is an indecomposable module and that ϕ is a non-invertible endomorphism of M . Let N be the image of ϕ , which is necessarily a proper submodule of M and a proper quotient of M . This implies that there are non-zero morphisms in both directions between M and N , contrary to our assumption. Therefore all indecomposable modules are bricks.

If τM is non-zero, there is a short exact sequence

$$0 \rightarrow \tau M \rightarrow E \rightarrow M \rightarrow 0.$$

Thus, if $\text{Hom}(M, \tau M) \neq 0$, then there is a cycle in $\text{mod } A$. Thus, all indecomposable objects are also τ -rigid. \square

We specialize further by imposing a finiteness assumption. We assume that A is representation-finite, i.e., it has only finitely many indecomposable representations up to isomorphism. By Proposition 7.1, we could just as well have assumed only that the number of τ -rigid indecomposable modules is finite.

7.2. Torsion classes

A *torsion class* in $\text{mod } A$ is a full additive subcategory of $\text{mod } A$ closed under extensions and quotients. We write $\text{tors } A$ for the torsion classes in $\text{mod } A$. Torsion classes naturally form a lattice, since the intersection of two torsion classes is again a torsion class. Since A is representation-finite, there are only finitely many different torsion classes, so $\text{tors } A$ is a finite lattice. We already showed in [TW17] that this lattice is trim.

There is a dual notion to torsion classes, the torsion-free classes. A *torsion-free class* is a full additive subcategory closed under extensions and submodules. There is a natural inclusion-reversing correspondence between torsion classes and torsion-free classes: if \mathcal{T} is a torsion class, then

$$\mathcal{F} = \mathcal{T}^\perp = \{Y \mid \text{Hom}(X, Y) = 0 \text{ for all } X \in \mathcal{T}\}$$

is the corresponding torsion-free class. The set of indecomposable objects of \mathcal{F} is maximal with respect to the property of having no morphisms from a module in \mathcal{T} (or equivalently, from an indecomposable module in \mathcal{T}). The dual statement is also true: the set of indecomposable modules in \mathcal{T} is maximal with respect to the property of having no nonzero morphisms into a

module in \mathcal{F} (or equivalently, into an indecomposable module in \mathcal{F}). Further, every pair $(\mathcal{T}, \mathcal{F})$ which is maximal in both these respects is automatically a *torsion pair*—a torsion class and its corresponding torsion-free class. See [ASS06, Section VI.1].

In this way, we see the connection with maximal orthogonal pairs. Consider the graph G whose vertices are indexed by the indecomposable A -modules, with an arrow from M to N if and only if $\text{Hom}(M, N) \neq 0$. By our assumption that there are no cycles in $\text{mod } A$, the graph G is acyclic. It is now clear, as we observed in [TW17], that maximal orthogonal pairs of G correspond to torsion pairs.

Theorem 7.2 ([TW17, Corollary 1.5]). *If A is representation finite and $\text{mod } A$ has no cycles, then the maximal orthogonal pairs in the trim lattice $\mathcal{L}(G)$ are naturally the torsion pairs of A , ordered with respect to inclusion of torsion classes (or reverse inclusion of torsion-free classes).*

7.3. Semibricks, flips, and the edge-labelling of the Hasse diagram of the lattice of torsion classes

There is a natural labelling of each edge of the Hasse diagram of the lattice of torsion classes of $\text{mod } A$ by an indecomposable module—if $\mathcal{U} < \mathcal{V}$ is a cover relation in the lattice of torsion classes, then its label is the unique brick in $\mathcal{V} \cap \mathcal{U}^\perp$ [Asa16, Proposition 1.17]. Comparing with our Theorem 4.2, we conclude that the labelling of cover relations by bricks coincides with the labelling coming from the overlapping cover relations of mops.

A collection of bricks is called a *semibrick* if there are no morphisms between two non-isomorphic bricks in the collection. Asai showed that there is a bijection between torsion classes and semibricks [Asa16, Theorem 1.3]. The semibrick $s(\mathcal{T})$ corresponding to a torsion class \mathcal{T} is the collection of labels on edges down from \mathcal{T} [Asa16, Lemma 1.16, Proposition 1.17]. It can also be described as the unique semibrick such that \mathcal{T} consists of modules filtered by quotients of elements of $s(\mathcal{T})$ [Asa16, Lemma 1.5]. Dually, there is a semibrick corresponding to each torsion free class: the semibrick $s(\mathcal{T}^\perp)$ corresponding to the torsion free class \mathcal{T}^\perp is the collection of labels on the edges up from \mathcal{T} . This semibrick is the unique semibrick such that \mathcal{T}^\perp is filtered by submodules of the elements of $s(\mathcal{T}^\perp)$.

An analogue of Theorem 4.9 allows us to directly compute the semibrick corresponding to a given torsion class.

Proposition 7.3. *The semibrick associated to a torsion class \mathcal{T} can be obtained by greedily building a semibrick by adding modules from \mathcal{T} if possible in any linear extension of the G -order.*

Proof. Let $\mathcal{D} = \{X_1, \dots, X_r\}$ be the semibrick associated to \mathcal{T} . As we consider modules from \mathcal{T} in some linear order compatible with G , suppose that we have already added $\{X_1, \dots, X_i\}$ and we encounter some module Y which does not have any morphisms from X_1, \dots, X_i . Since Y is in \mathcal{T} , it is filtered by quotients of modules in \mathcal{D} . Now Y does not have any morphisms from the already-chosen elements of \mathcal{D} , but Y will not admit any morphisms from any subsequent module. Thus Y must be an element of \mathcal{D} as well. On the other hand, if Y admits a morphism from some module which we have already added to \mathcal{D} , then obviously we must not add it into \mathcal{D} , and following our procedure, we do not. We will therefore successively add all the elements of \mathcal{D} by following this procedure. \square

Using semibricks, we can provide a representation-theoretic justification for our definition of flips on tight orthogonal pairs in Algorithm 3.

Theorem 7.4. *Let $(\mathcal{D}, \mathcal{U})$ be the tight orthogonal pair associated to the torsion class \mathcal{T} . Let $X \in \mathcal{U}$. Let $(\mathcal{D}', \mathcal{U}')$ be the tight orthogonal pair associated to the torsion class \mathcal{T}' which covers \mathcal{T} along an edge of the Hasse diagram labelled by X . Then Algorithm 3 successfully reconstructs $(\mathcal{D}', \mathcal{U}')$.*

Proof. The torsion class \mathcal{T}' is the minimal torsion class containing X and \mathcal{T} . Clearly it contains all modules which are filtered by quotients of X and elements of \mathcal{T} , and since that category is extension closed and quotient closed, it must be \mathcal{T}' . In particular, we observe that when restricted to modules that do not follow X in G -order, the two classes \mathcal{T}' and \mathcal{T} coincide, which we write as $\mathcal{T}'|_{\not\geq X} = \mathcal{T}|_{\not\geq X}$. Therefore, by Proposition 7.3, $\mathcal{D}|_{\not\geq X} = \mathcal{D}'|_{\not\geq X}$. Now

$$\mathcal{F}' = \mathcal{T}'^\perp = \{Y \in \mathcal{F} \mid \text{Hom}(X, Y) = 0\}.$$

Note that $\mathcal{U}|_{\not\geq X} \subseteq \mathcal{F}'$, since if we reconstruct \mathcal{U} using (the dual version of) Proposition 7.3, we must add in X , and therefore we must not have added any module which admits a map from X previously. Now, since $\mathcal{U}|_{\not\geq X}$ is contained in $\mathcal{F}' \subset \mathcal{F}$, we see that $\mathcal{U}'|_{\not\geq X} = \mathcal{U}|_{\not\geq X}$.

The remaining elements of \mathcal{D}' , i.e., those which follow X in G -order, can then be reconstructed by Proposition 7.3. We first observe that X itself is in \mathcal{D}' , since it is in \mathcal{T}' , and since $X \in \mathcal{U}$, there are no morphisms from $\mathcal{D}'|_{\not\geq X}$ to X . To use Proposition 7.3 to determine the further elements of \mathcal{D}' , we first need to know the elements of $\mathcal{T}'|_{>X}$. These are the modules strictly following X which do not admit a morphism into \mathcal{F}' . We can replace \mathcal{F}' by \mathcal{U}' , and then we can replace \mathcal{U}' by $\mathcal{U}'|_{>X}$, since modules following X only have morphisms into other modules following X . But we have already observed that $\mathcal{U}'|_{>X} = \mathcal{U}|_{>X}$.

We therefore see that the modules of \mathcal{D}' which strictly follow X are constructed greedily in G -order maintaining the properties of having no morphisms from any of the other modules in \mathcal{D}' , and having no morphisms to \mathcal{U} .

This construction of \mathcal{D}' is manifestly the same as the construction given in Algorithm 3, as desired. The argument for the construction of \mathcal{U}' is essentially dual. □

7.4. 2-simple-minded collections

A *simple-minded collection* for A is a collection of objects X_1, \dots, X_r in the derived category of $D^b(A)$ such that:

1. $\text{Hom}(X_i, X_j[m]) = 0$ for $m < 0$,
2. $\text{End}(X_i)$ is a division algebra and $\text{Hom}(X_i, X_j) = 0$ unless $i = j$,
3. X_1, \dots, X_r generate $D^b(A)$ in the sense that the smallest thick subcategory containing all of them is $D^b(A)$ itself,

See [KY14] for more on simple-minded collections. A *2-simple-minded collection* has the additional property that for each X_i , we have $H^j(X_i) = 0$ for $j \neq 0, -1$. It turns out that the elements of a 2-simple-minded collection are all contained in $\text{mod } A \cup \text{mod } A[1]$, see [BY14, Remark 4.11].

Asai showed that there is a bijection from torsion classes to 2-simple-minded collections [Asa16, Theorem 2.3], which sends \mathcal{T} to $s(\mathcal{T}) \cup s(\mathcal{T}^\perp)[1]$.

It follows from our bijection between torsion classes and mops, and the labelling of the cover relations by bricks and Theorem 4.2, that tight orthogonal pairs are in bijection with 2-simple-minded collections.

Theorem 7.5. *Let A be a representation finite k -algebra with no cycles in $\text{mod } A$. There is a bijection from tight orthogonal pairs to 2-simple-minded collections, sending $(\mathcal{D}, \mathcal{U})$ to $\mathcal{D} \cup \mathcal{U}[1]$.*

Although Theorem 7.5 follows from our results together with those of [Asa16], we find it instructive to give a direct proof of the following proposition in order to show how the tightness condition on tops naturally arises from 2-simple-minded collections.

Proposition 7.6. *If $\mathcal{D} \cup \mathcal{U}[1]$ is a 2-simple-minded collection, then $(\mathcal{D}, \mathcal{U})$ is a tight orthogonal pair.*

Proof. Let $\mathcal{C} = \{X_1, \dots, X_r, Y_1[1], \dots, Y_s[1]\}$ be a 2-term simple minded collection. From the condition (2), we see that $\mathcal{D} = \{X_1, \dots, X_r\}$ and $\mathcal{U} = \{Y_1, \dots, Y_s\}$ are independent sets. The fact that $\text{Ext}^{-1}(X_i, Y_i[1]) = 0$ implies that $\text{Hom}(X_i, Y_j) = 0$, i.e., that $(\mathcal{D}, \mathcal{U})$ is orthogonal.

The proof of tightness is somewhat more involved. Write \mathcal{T} for the torsion class associated to \mathcal{D} , and \mathcal{T}^\perp for the torsion-free class associated to \mathcal{U} . If Z is a module in \mathcal{T} , then it cannot be added into \mathcal{D} because such a module is filtered by quotients of objects from \mathcal{D} , and thus admits a morphism from one of them, which is forbidden. If Z is a module not in \mathcal{T} , then it admits a map to a non-zero object in \mathcal{T}^\perp , and thus to an object in \mathcal{U} ; this is also forbidden.

It is also impossible to replace X_i by X'_i which is above X_i . If X'_i is not in \mathcal{T} , then this is impossible for the reason given above that we cannot add X'_i into \mathcal{D} .

On the other hand, if X'_i is in \mathcal{T} , then it is filtered by quotients of elements of \mathcal{D} , so it admits a morphism from some element of \mathcal{D} , and since X'_i is strictly above X_i , this element cannot be X_i . Thus X_i cannot be replaced by X'_i in this case either.

Dual considerations explain why no element can be added to \mathcal{U} and no element of \mathcal{U} can be lowered. \square

Putting together the previous results, we conclude the following theorem.

Theorem 1.9. *If A is representation finite and $\text{mod } A$ has no cycles, then maximal orthogonal pairs correspond to torsion pairs, independent sets correspond to semibricks, and tight orthogonal pairs correspond to 2-simple-minded collections.*

Acknowledgements

We thank two diligent anonymous referees for helping us improve our exposition.

References

- [Asa16] Sota Asai, *Semibricks*, arXiv preprint arXiv:[1610.05860](#) (2016).
- [ASS06] Ibrahim Assem, Daniel Simson, and Andrzej Skowroński, *Elements of the representation theory of associative algebras, volume 1: Techniques of representation theory*, London Mathematical Society Student Texts, vol. 65, Cambridge University Press, 2006. [MR2197389](#)
- [BY14] Thomas Brüstle and Dong Yang, *Ordered exchange graphs*, Advances in Representation Theory of Algebras (2014), 135–193. [MR3220536](#)

- [CF95] Peter J Cameron and Dmitry G Fon-der-Flaass, *Orbits of antichains revisited*, European Journal of Combinatorics **16** (1995), no. 6, 545–554. [MR1356845](#)
- [Fon93] Dmitry G Fon-der-Flaass, *Orbits of antichains in ranked posets*, European Journal of Combinatorics **14** (1993), no. 1, 17–22. [MR1197471](#)
- [JR17] Michael Joseph and Tom Roby, *Toggling independent sets of a path graph*, arXiv preprint arXiv:[1701.04956](#) (2017). [MR3678660](#)
- [KY14] Steffen Koenig and Dong Yang, *Silting objects, simple-minded collections, t-structures and co-t-structures for finite-dimensional algebras*, Doc. Math **19** (2014), 403–438. [MR3178243](#)
- [Mar75] George Markowsky, *The factorization and representation of lattices*, Transactions of the American Mathematical Society **203** (1975), 185–200. [MR0360386](#)
- [Mar92] George Markowsky, *Primes, irreducibles and extremal lattices*, Order **9** (1992), no. 3, 265–290. [MR1211380](#)
- [Str16] Jessica Striker, *Rowmotion and generalized toggle groups*, arXiv preprint arXiv:[1601.03710](#) (2016). [MR3811480](#)
- [SW12] Jessica Striker and Nathan Williams, *Promotion and rowmotion*, European Journal of Combinatorics **33** (2012), no. 8, 1919–1942. [MR2950491](#)
- [Tho06] Hugh Thomas, *An analogue of distributivity for ungraded lattices*, Order **23** (2006), no. 2, 249–269. [MR2308910](#)
- [TW17] Hugh Thomas and Nathan Williams, *Rowmotion in slow motion*, arXiv preprint arXiv:[1712.10123](#) (2017).

HUGH THOMAS
LACIM
UNIVERSITÉ DU QUÉBEC À MONTRÉAL
CANADA
E-mail address: hugh.ross.thomas@gmail.com

NATHAN WILLIAMS
UNIVERSITY OF TEXAS AT DALLAS
USA
E-mail address: nathan.f.williams@gmail.com

RECEIVED 21 APRIL 2018